Intensity-Based Edge Classification

Andrew P. Witkin

Fairchild Laboratory for Artificial Intelligence Research
Palo Alto, CA 94304

ABSTRACT

A new intensity-based approach to the classification of edges was developed and implemented. Using basic continuity and independence properties of scenes and images, signatures were deduced for each of several edge types expressed in terms of correlational properties of the image intensities in the neighborhood of the edge. This procedure's ability to discriminate occluding contours from cast shadow boundaries was demonstrated for cases where line junction cues are absent from the image.

1. Introduction

Edges in images arise from several very different kinds of scene event — occluding contours, discontinuities of surface orientation, material changes, cast shadows, etc. Each kind of edge contributes its own significant constraints to image interpretation[1-8] but these constraints cannot be exploited unless edge types can be distinguished from each other. Edge classification is therefore of considerable importance to image interpretation.

Edge classification in idealized line drawings has been treated in terms of junction constraints,[1-3] but the perfect line drawings required to identify junctions have proved difficult to obtain from natural images. Horn[9] has suggested that the intensity profiles across edges (peak vs. step, etc.) may provide distinguishing signatures for some edge types. However, this technique depends on quantitative photometry, and its effectiveness has never been demonstrated for complex imagery.

This paper describes a classification technique that, like line-junction methods, relies on structural rather than quantitative photometric properties of the image and scene, but unlike those methods, utilizes the raw image intensities in the neighborhood of the edge, without requiring elaborate analyses of edge structure. The method follows from basic properties of image edges — occluding contours and cast shadows in particular will be considered —and from basic properties of scene structure. Occluding contours are "seams" in the projective fabric, curves across which surface points that may be widely separated in space are juxtaposed in the image by the vagaries of projection. Cast shadows are curves across which the shadowed surface's image undergoes a systematic (ideally, linear) transformation. Given this characterization, two basic complementary properties of scenes suffice to distinguish cast shadows and occluding contours from each other and from arbitrary image curves: (1) the processes that shape, color, and illuminate natural surfaces act continuously almost everywhere, and (2) the properties of widely separated scene constituents are independent.

Given these properties, we conclude that intensities at nearby points on either side of an arbitrary curve in the image are likely to be highly correlated, due to the coherence of surface structure; intensities at nearby image points across an occluding contour are likely to be uncorrelated, because they are actually the projections of distant, hence independent, surface points; and intensities across a cast shadow edge are likely to be highly correlated under a systematic (ideally, linear) transformation, arising from the interaction of the continuous underlying surface structure with the illumination transition.

These observations are applied to image edges by constructing a family of parallel curves around the edge, and performing a sequence of linear regressions of the intensity values along each curve onto those along its neighbor. The following behavior is predicted for each of several edge types:

- A precipitous drop in correlation at the nominal edge location signifies an occluding contour.
- High correlation with an abrupt shift in the regression parameters signifies a shadow.
- Sustained high correlation with the additive and multiplicative regression parameters near zero and unity, respectively, implies that no significant edge is present.
- Low correlation throughout implies lack of coherent surface structure, and no edge type can be assigned.

These ideal structural "signatures" were shown to closely predict regression sequences obtained from images of natural edges.

2. Defining the Problem

Because edge types are defined in terms of the scene events they denote, any method for edge sorting must provide some basis for discriminating those events by their appearance in the image. We therefore begin by characterizing the distinctive properties of occluding contours and cast shadow edges, and defining the computational problem of identifying those edges.

Occluding contours: the projective mapping from image to scene tends to be continuous, because physical surfaces tend to be continuous. Almost everywhere in a typical image, therefore, nearby points in the image correspond to nearby points in the scene. This adjacency is preserved over any change in point of view or scene configuration, short of rending the connected surfaces of which the scene is composed. The distinguishing property of occluding contours (which may be defined as discontinuities in the projective mapping) is their systematic *violation* of this rule: the apparent juxtaposition of two surfaces across an occluding edge represents no fixed property of either surface, but is subject to the vagaries of viewpoint and scene configuration. For example, if you position your finger to coincide with a particular feature on the wall or outside the window, a small change in the position of head or hand may drastically affect their apparent relation. Because the false appearance of proximity is the hallmark of occluding edges, the problem in identifying those edges may be cast as that of distinguishing in the image the actual proximity of nearby points on connected surfaces from accidental proximity imposed by projection.

Cast shadows: cast shadows in outdoor scenes usually represent transitions from direct to scattered illumination caused by the interposition of an occluding body between the sun and the viewed surface. The problem in identifying cast shadows is to distinguish these transitions in incident illumination from changes in albedo, surface orientation, and so forth. This discrimination presents a problem because the effects of all these parameters are confounded in the image data— a change in image brightness may reflect a change in albedo or surface orientation, as well as incident illumination. Because the relation among illumination, reflectivity, orientation, and image irradiance is well known, the presence of shadows in an image could be readily detected if a constant reference pattern could be placed in the scene: when the apparent brightness of a constant pattern varies with location, the change in brightness must, by elimination, be attributed to a change in illumination. Of course such active intervention is generally impractical. The problem may be viewed as that of achieving the effect of viewing a constant pattern across the shadow edge, without actually placing such a pattern in the scene. This could be achieved if some fixed relation were known to hold between the surface strips on each side of the shadow edge.

In short, occluding contours are curves across which points that may be distant in space are placed in apparent juxtaposition by projection, violating the continuity of the projective mapping that holds over most of the image. To identify occluding contours therefore requires that actual proximity be distinguished from apparent proximity imposed by projection. Cast shadows edges are contours across which the pattern of surface reflectance has been systematically transformed by an abrupt change in illumination. To identify cast shadow edges, the effects of illumination must be distinguished from those of albedo and surface orientation, as if a constant reference pattern had been placed across the edge.

3. Continuity and independence

The solution we have devised rests on two simple complementary principles: (1) **Continuity:** Surfaces, surface markings, and illumination are almost everywhere continuous. Therefore, the projective mapping is almost everywhere continuous, and image intensities at nearby points tend to be highly correlated. (2) **Independence:** The factors governing the structure of a scene —the shapes of objects, and their placement with respect to each other, to illuminants, and to the viewer —are so complex that properties of distinct or widely separated scene constituents may for most purposes be regarded as causally independent. (This independence principle is related to the principle of general position,[4],[8] which assumes isotropy for viewpoint and object position and orientation.)

One simple measure of continuity or coherence across an image curve (there are many others) is linear correlation between the image intensities a small distance on either side of the curve: a high positive correlation implies that the image strips on either side of the curve are closely related, a low correlation implies no (linear) relation. Given a high correlation, a regression equation can be computed to describe the linear transform relating the intensities across the curve. Several predictions about correlations and regressions across edges follow from the continuity and independence principles:

The continuity principle implies that high correlations should often be observed across arbitrarily selected curves in the image. However, a low correlation could just imply low contrast or fragmented surface structure.

The independence principle implies that high correlations should almost *never* be observed across occluding contours, because the points meeting along those curves are not the projections of nearby points in space.

The independence principle implies that a cast shadow edge would not have any unusual properties, were the shadowing body removed, because the light source and the shadowed and shadowing objects do not "conspire" to acheive special alignments. Therefore, apart from the effects of the shadow itself, a cast shadow edge should show the same correlational properties as an arbitrarily selected curve. Shadows in outdoor scenes are often transitions between two roughly constant levels of illumination —scattered and direct —and the effect of an illuminant change on image intensity may be very roughly idealized as linear. To the extent these idealizations hold, linear correlations across cast shadow edges are likely to be as high as those across arbitrary curves, but the illumination transition will appear as a perturbation of the regression equation (ideally, a multiplictive factor for linear digitization, an additive one for logarithmic digitization.)

Given a candidate edge in the image, these observations leave us with the strong implication that a high correlation of intensities across the edge *excludes* the possibility that the edge is an occluding contour, and the somewhat weaker implications that a low correlation signals an occluding contour, and that a high correlation through a substantial linear transform signals a shadow. The latter implications are weaker because lack of correlation may just signify conditions, such as low contrast, that don't favor correlation. Our conclusions about occluding contours and shadows may be strengthened by examining a larger neighborhood around the edge. By embedding the given edge in a series of parallel curves, a sequence of regressions can be performed one onto the next. A low correlation throughout signals low contrast or lack of texture, and no conclusion can be drawn. However, a sharp notch in an otherwise high correlation, where the regression sequence crosses the edge, argues against global low contrast or lack of texture, providing a stronger indicator that the edge is an occluding contour. Likewise, sustained high correlation with an abrupt perturbation in one or both regression paramaters is good evidence for a shadow. Finally, sustained high correlation with no perturbation of the regression equation provides evidence that the edge is not physically significant.

4. Implementation and results

Our implementation assumes that an edge has been located by edge-finding techniques. Hand-traced edges and zero-crossings in a $\nabla^2 G$ convolution[10] were tried as inputs. A parallel family of curves was constructed around the edge as follows: at fixed intervals of arc length along the edge, a line normal to the edge was constructed. The set of points lying some fixed distance from the edge along each normal line then defines a "parallel" curve. A family of curves was constructed by varying that perpendicular distance. This construction amounts to warping a strip of the image, surrounding the edge, into a rectangular region, whose central column corresponds to the original edge. The vertical dimension of this "rectified strip" denotes arc length on the edge, and the horizontal dimension denotes perpendicular distance from the edge. The columns surrounding the central one correspond to parallel curves on either side of the edge. Intensity values for the rectified strip were obtained by bilinear interpolation of the intensities in the original image, to reduce quantization error.

Once the rectified strip was constructed, a sequence of linear regressions was performed between columns. To avoid spurious correlation imposed by the imaging and digitizing process, regressions were computed between the ith column and the $(i+2)$th. The outcome of this computation was a normalized correlation, an additive regression term, and a multiplicative regression term, each a function of column position. The midpoints of these plots represent the regression across the original edge. Points to either side of the plots' midpoints represent regressions between adjacent parallel curves on either side of the original edge.

The idealized edge type "signatures", as developed in the preceding section, are shown and explained in fig. 1, in terms of these regression sequence plots.

Actual image edges, together with rectified strips and regression plots, are shown and described in fig.s 2-5. (The edges in these cases were hand-traced.) Marked correspondence between the actual and idealized plots is evident to inspection. No attempt has yet been made to automate this comparison, although a variety of simple thresholding schemes might well prove empirically adequate to the classification task.

5. Conclusions

Two conclusions may be drawn from these results: first, the correlational properties of intensities in the neighborhood of an edge carry important information about the edge's physical significance, and second, the very basic principles of continuity and independence can provide strong an useful constraints on image interpretation.

Acknowledgements

This work was performed while the author was a member of the Artifical Intelligence Center at SRI, International. Support was provided in part by NASA contract #NAS1-16282, NSF grant #MCS-7901830, and DARPA contract #DAAG-29-79-C-0216.

REFERENCES

[1] Huffman, D.A, Impossible objects as nonsense sentences, in:, Machine Intelligence, Edinburg University Press, Edinburgh, 1971.

[2] Clowes, "On seeing things", *Artificial Intelligence*, **2**(1971), 79—112.

[3] Waltz, D.L., Generating semantic descriptions from drawings of scenes with shadows, AI-TR-271, MIT, Cambridge, MA, 1972.

[4] Marr, D. C., "Analysis of occluding contour", *Proc. Roy. Soc. Lond.*, **197**(1977), 441—475.

[5] Stevens, K.A., "The visual interpretation of surface contours", *Artificial Intelligence*, **17**(1981), 47—74.

[6] Witkin, A.P, "Recovering surface shape and orientation from texture", *Artificial Intelligence*, **17**(1981), 17—46.

[7] Barrow, H.G. and Tenenbaum, J.M., "Interpreting line drawings as three dimensional surfaces", *Artificial Intelligence*, **17**(1981), 75—116.

[8] Binford, T.O, "Inferring surfaces from images", *Artificial Intelligence*, **17**(1981), 205—244.

[9] Horn, B. K. P., "Understanding Image Intensities", *Artificial Intelligence*, **21**, 11(1977), 201—231.

[10] Marr, D. C. & Hildreth, E., "A theory of edge detection", *MIT AI Memo 518*(1979).

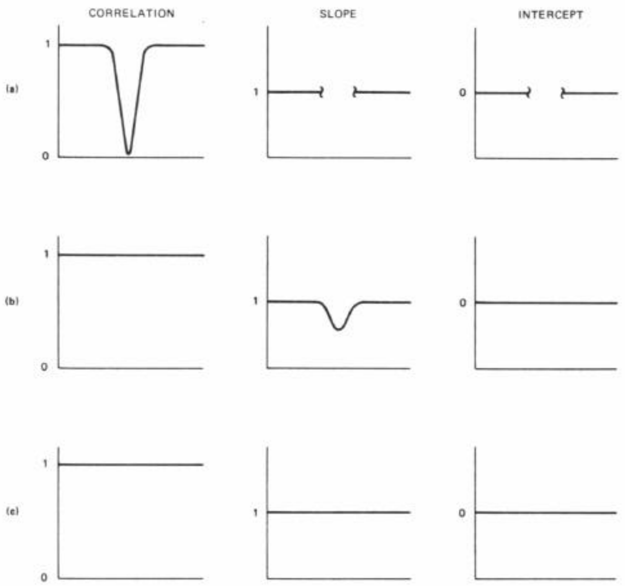

FIGURE 1—Idealized edge signatures. Each row of plots represents an idealized sequence of linear regressions across an edge. **(a) Occluding contour:** a sharp notch in correlation across the edge. The plots for the slope and intercept of the regression equation are broken to indicate that the regression equation is meaningless in the low-correlation area. **(b) Cast shadow:** Sustained high correlation across the edge, with disturbance of one or both regression parameters. The nature of this disturbance depends on the sense of the edge (i.e. whether the shadow lies on the left or right), and on details of the imaging and digitizing process. In practice, nonlinearities perturb the correlation slightly. **(c) No edge present:** sustained high correlation without disturbance in the regression parameters implies that the edge is not physically significant. An additional case, not illustrated, is that of low correlation throughout. This indicates low contrast or lack of surface structure, implying that no descision about edge type can be made.

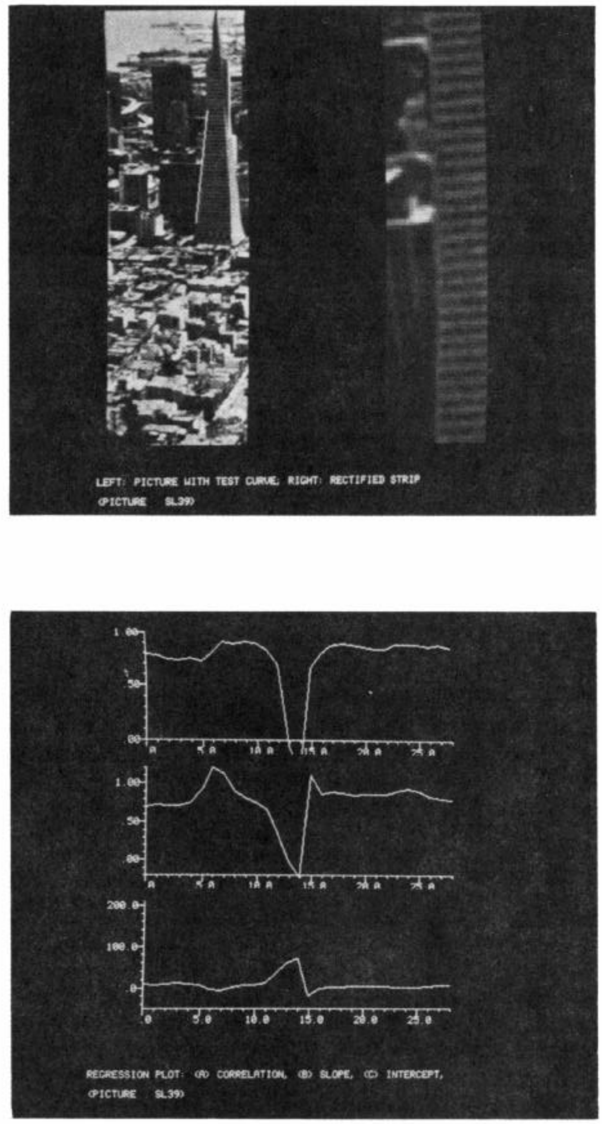

FIGURE 2—Example of an occluding edge. In the upper left is the original image, with the selected edge highlighted. The upper right shows the rectified strip, whose midline corresponds to the edge. Below are the plots for correlation, slope, and intercept. The overall high correlation, with a sharp plunge near the edge location, corresponds closely to the idealized form of Fig. 1a. Remember that slope and intercept of the regression line are meaningless where the correlation is low.

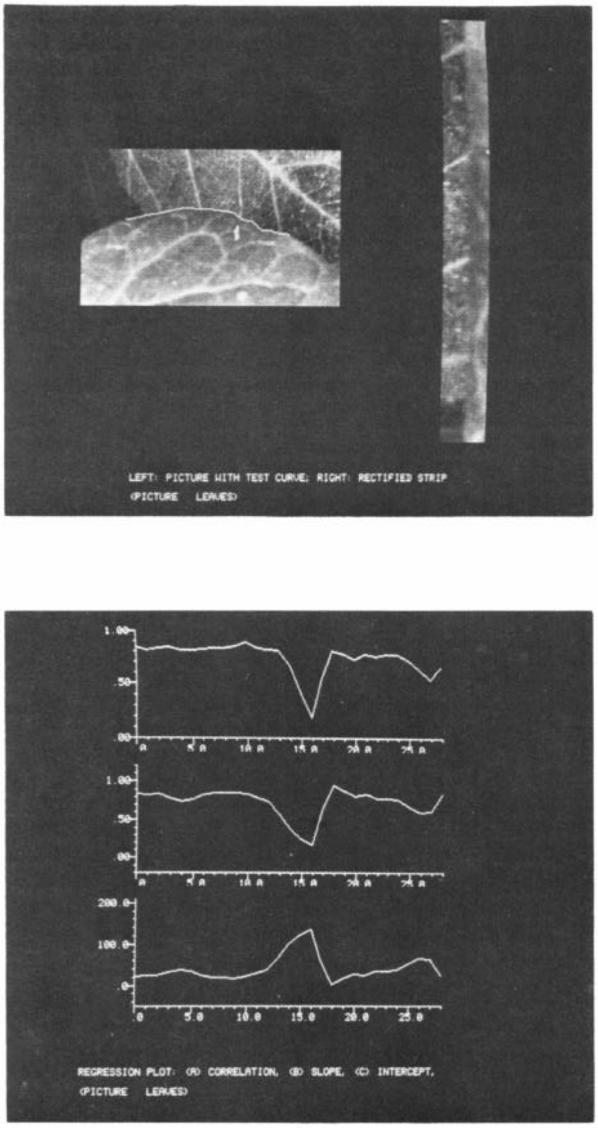

FIGURE 3—A low-contrast occluding edge. Although the edge contrast is low, the correlation across the edge dips to about .15, compared to about .85 in the surrounding region.

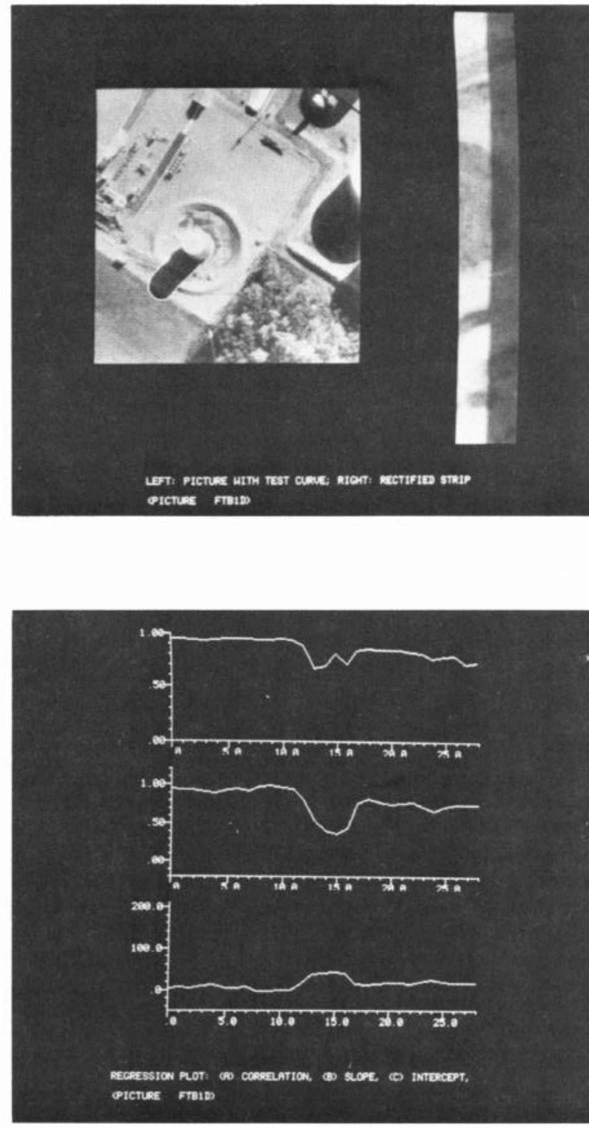

FIGURE 4—A cast-shadow edge. The shadow transition appears primarily as a dip in the slope of the regression line (a dip rather than a bump because the left-to-right transition is from light to shadow.) The perturbation in the additive regression term, and the dip in correlation to about .75, are due primarily to nonlinearities in the film response.

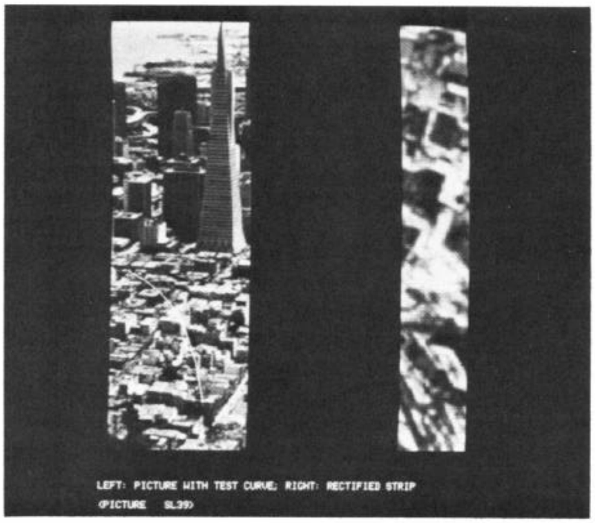

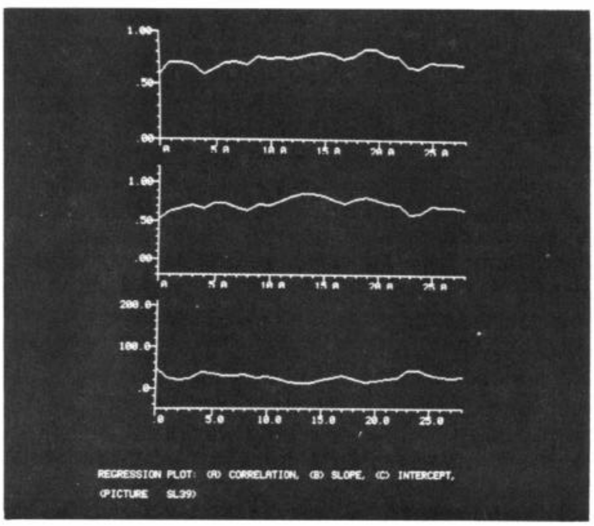

FIGURE 5— Regression where no physical edge is present. As expected, a fairly high correlation, with stable regression parameters, is maintained across the "edge".