The Effect of Spatial Discretization on the Magnitude and Direction Response of Simple Differential Edge Operators on a Step Edge Part1: Square Pixel Receptive Fields

L.J. Kitchen and J.A. Malin

COINS Technical Report 87-34

April 1987

This research has been supported under grants: NSF DCR-8318776, AFOSR-86-0021, DMA-800-85-C-0012, and ARMY ETL DACA76-85-C-0008.

The Effect of Spatial Discretization on the Magnitude and Direction Response of Simple Differential Edge Operators on a Step Edge Part 1: Square Pixel Receptive Fields

L. J. Kitchen
J. A. Malin *
Computer and Information Science Department
University of Massachusetts
Amherst, MA 01003

April 10, 1987

Abstract

For a unit step edge, we calculate the gradient magnitude and direction reported by various simple differential edge operators as a function of the edge's actual orientation and offset with respect to the pixel grid. To be consistent with previous work (of Abdou and Pratt) we have initially chosen to analyse (among others) the Sobel, Prewitt, and Roberts Cross operators, using a digitization model in which the pixels have uniform square receptive fields that tesselate the image plane. Our quantitative results provide insights into the behavior of these commonly used operators, insights that can guide their proper application in problems of image understanding and robot vision. We also suggest novel techniques for improving the performance of these operators.

*This research has been supported under grants: NSF DCR-8318776, AFOSR-86-0021, DMA-800-85-C-0012, and ARMY ETL DACA76-85-C-0008

1 Introduction

Despite the development of more sophisticated edge detectors [15,10,5,11,17], simple operators such as the Sobel [7], Prewitt [21], and Roberts Cross [22] are well established and commonly used, chiefly because of their computational simplicity, a consideration of particular importance in applying computer vision to practical tasks such as robotics. Therefore it is important to study the behavior of these operators, in order to discover their strengths and weaknesses, so as to make best, appropriate use of them in applications.

We are particulary interested in the extraction of straight-edge features for visual guidance of autonomous robot vehicles, and for recognition of polyhedral objects using techniques based on that of Burns [4]. For this task, a detailed understanding of edge-operator performance is crucial.

Because of the importance of edge operators, there have been numerous comparative studies evaluating the efficacy of various edge operators [8,9,6,2,3,14]. These studies generally are of two types: analytical or empirical. In analytical studies, some formal model of an edge is used and the behavior of the operator on such an edge is analysed mathematically in terms of the parameters of the edge and the operator. In empirical studies, the edge operator is applied to images, either natural or synthetic, and statistics of its performance gathered. These statistics may be on the fidelity of the operator to known edge locations in the image [2], [8], or on some measure of intrinsic edge coherence [14]. Of course, there is no sharp distinction between the analytical and the empirical. An empirical study may use a precise mathematical edge model in generating synthetic images; an analytical study may use numeric simulation in obtaining its results. Also, authors may use combinations of these two approaches in a single study [2].

The study described here is basically an analytical one. It follows on most immediately from some work of Abdou and Pratt. We use the same step-edge model, and include all "differential" edge operators that they look at. However, in their study Abdou and Pratt present results for all significant edge orientations only for one special offset with respect to the pixel grid, and for a limited range of offsets for only two special orientations. Our study goes beyond theirs in that we consider fully the joint effect of edge orientation and offset over the full range of significant offsets; the study thus elucidates more fully the behavior of these operators over the entire variation in the edge model. We also propose simple techniques for improving the performance of these operators.

2 The Edge and Digitization Model

The edge model we use is of an infinite straight step edge. Since all the operators we examine are linear in the image intensity, we can without loss of generality take the edge to be a step edge of unit contrast; zero intensity on one side, unit intensity on the other. Thus the edge is characterized only by the (directed) straight line defining the boundary. The most convenient and natural parameterization of this line is in terms of its orientation θ and perpendicular offset ρ from the origin, with θ chosen so that it points from the dark side of the edge to the bright. See Figure 1. We choose this model chiefly for consistency with Abdou and Pratt, but it is a reasonable model for the local appearance of many edges encountered in practice.

Our digitization model is this: Imagine the image plane tessellated with a square grid. Without loss of generality we can take the sides of these squares to be of unit length. A pixel is identified with each of these squares, and takes this square area as its receptive field. That is, each pixel integrates the light intensity over its associated square, with uniform weighting over the square.

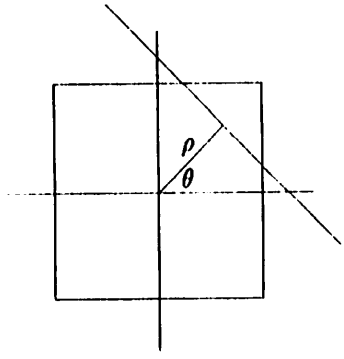

Figure 1: Parameterization of straight step edge passing through a pixel.

Digitization is more fully treated in [20].

So, for a given pixel and edge, the intensity for that pixel can be computed as follows: First, because of the symmetry of the pixel grid and of the operators examined, we need consider only orientations in the range zero to 45°. Take the origin to be at the center of the pixel. Then for an edge with orientation θ and offset ρ, the integrated intensity $I(\rho,\theta)$ for the pixel is merely the area of the intersection of the pixel with the bright side of the edge. It is given by:

$$I(\rho,\theta) = \begin{cases} 1 - I(-\rho,\theta) & \text{for } \rho < 0 \\ 0 & \text{for } \theta = 0°,\ \rho \geq \frac{1}{2} \\ \frac{1}{2} - \rho & \text{for } \theta = 0°,\ 0 \leq \rho < \frac{1}{2} \\ 0 & \text{for } \theta > 0°,\ \frac{1}{2} \leq \eta \\ \frac{\tan\theta}{8} + \frac{1}{4} - \frac{\rho}{2\cos\theta} + \frac{1}{8\tan\theta} - \frac{\rho}{2\sin\theta} + \frac{\rho^2}{\sin 2\theta} & \text{for } \theta > 0°,\ -\frac{1}{2} \leq \eta < \frac{1}{2} \\ \frac{1}{2} - \frac{\rho}{\cos\theta} & \text{for } \theta > 0°,\ \eta < -\frac{1}{2} \end{cases}$$

where $\eta = \rho\csc\theta - 1/2\cot\theta$. The point $(1/2,\eta)$ is the intersection between the edge line and the right border line of the pixel square. This is a general formulation of the edge digitization model used by Abdou and Pratt for the three special cases they treated, as will be described below.

3 Edge Operators

All the operators considered here are *differential* operators (in the sense of Abdou and Pratt). That is, they use convolutions to compute approximations to the X and Y components of the image intensity gradient. Typically, edge magnitude and direction are obtained by taking the L_2 (Euclidean) norm and direction of this gradient vector. That is

$$M_2 = \sqrt{I_x^2 + I_y^2}$$

$$\tan\hat{\theta} = I_y / I_x$$

where I_x and I_y are respectively the X and Y components of the intensity gradient, M_2 is the magnitude (computed according to the L_2 norm), and $\hat{\theta}$ is the gradient direction, chosen to be in the proper quadrant. However, for computational reasons the L_1 and L_∞ norms are sometimes used instead, for the edge magnitude, respectively

$$M_1 = |I_x| + |I_y|$$

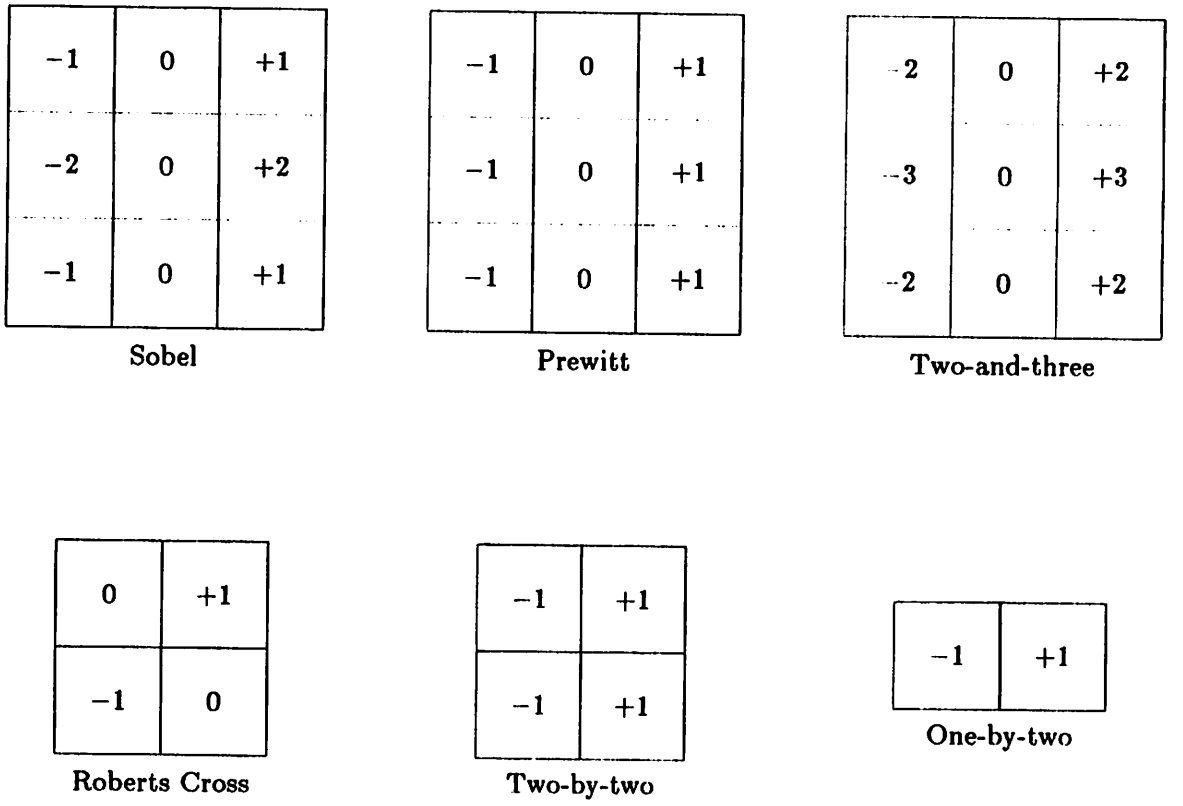

Figure 2: Convolution masks for I_x.

$$M_\infty = \max(|I_x|, |I_y|)$$

The operators we analyse are the Prewitt [21], Sobel [7], Roberts Cross [22], "One by Two" [4], "Two by two" [4], and "Two and Three", a hybrid of the Prewitt and Sobel operators devised during this study. The convolution masks used by these operators in computing I_x are shown in Figure 2. The masks used for I_y are rotations of these by 90°. We place the origin at the center of the neighborhood; that is, at the center of the central pixel for the 3×3 operators, and midway between pixel centers for the 2×2. Thus for each neighboring pixel, its relative displacement (Δ_x, Δ_y) from the neighborhood center is shown in Figure 3. For an edge with parameters (ρ, θ) with respect to the origin of the neighborhood, the intensity at the neighbor with displacement (Δ_x, Δ_y) is given by

$$N(\Delta_x, \Delta_y, \rho, \theta) = I(\rho + \Delta_x \cos\theta + \Delta_y \sin\theta, \theta)$$

Thus, for example, for the Sobel operator

$$\begin{aligned} I_x(\rho,\theta) \;=\; & N(+1,+1,\rho,\theta) + 2N(+1,0,\rho,\theta) + N(+1,-1,\rho,\theta) \\ & - N(-1,+1,\rho,\theta) - 2N(-1,0,\rho,\theta) - N(-1,-1,\rho,\theta) \end{aligned}$$

From this and the corresponding expression for I_y, we can compute the magnitude response M and direction response $\hat{\theta}$ for the Sobel operator, as a function of a straight edge's actual offset ρ and

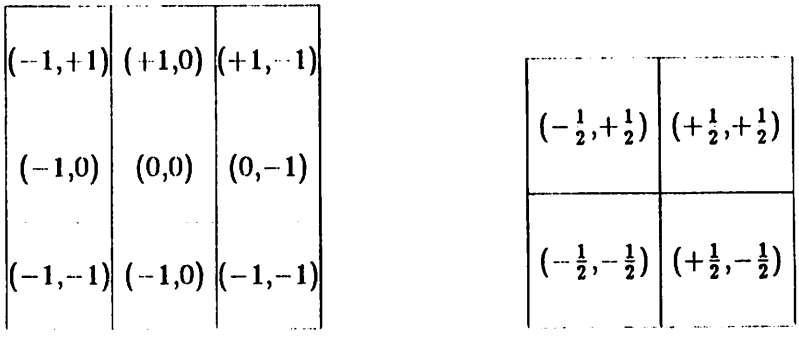

Figure 3: Relative displacements of neighbors from neighborhood center.

actual orientation θ. (Whether M is M_1, M_2, or M_∞ depends on the norm chosen.) The analogous computations can be done for the other operators.

Since the operators we consider are designed to compute the gradient of a smooth image surface, from which the digital image is obtained by sampling, it is hardly to be expected that they will report exact contrast and orientation for a discontinuous step edge. In a sense, it is inappropriate to apply these operators to a step edge. However, these operators *are* commonly used, and a step edge is a more realistic model for many image events than is a smooth surface, so it is certainly of interest to investigate how these operators perform on a step edge.

As mentioned earlier, by the symmetry of the pixel grid we need consider only edge orientations in the range 0° to 45°. To cover the responses to any edge passing through an operator's $n \times n$ neighborhood, we must consider offsets ρ that range from zero up to $(n/2)(\sin\theta + \cos\theta)$, which is $n/2$ at $\theta = 0°$, and $(n\sqrt{2})/2$ at $\theta = 45°$. Edges with parameters lying outside this region of (ρ, θ) space do not pass through the neighborhood.

However, in an image, if an edge passed entirely outside the unit square centered at the origin of the neighborhood, then that edge would be more strongly responded to at a different position in the pixel grid. Therefore the most interesting region of response is for edges that pass through this central unit square. Put another way, if among the discrete image positions at which the edge operator is applied we define an *edge position* to be a position such that an edge passes through the unit square centered at the origin of the operator's neighborhood, then we are particularly interested in the response at such edge positions. Of course, it is impossible to distinguish edge positions from non-edge positions locally, because the response for a neighborhood located directly on a weak edge will be indistinguishable from that for a neighborhood slightly offset from a stronger edge. However, the responses for edges passing through the central unit square are the best that the operator could be expected to do, even if we assume that there were some perfect method for determining edge positions.

Abdou and Pratt considered only three slices through the (ρ, θ) space, namely $0° \leq \theta \leq 45°$, $\rho = 0$; $\theta = 0°$, $0 \leq \rho \leq 1/2$; and $\theta = 45°$, $0 \leq \rho \leq \sqrt{2}/2$. While this gave some indication of an edge operator's response, it cannot characterize the behavior of an edge operator over the full variation of the edge model.

Operator	Minimum (-)	Maximum (+)	Avg (-)	Avg (+)
Prewitt	-7.429	1.837	-3.498	0.641
Sobel	-3.712	0.010	-1.066	0.010
Roberts Cross	-8.130	4.378	-4.446	1.402
One by Two	-8.060	15.982	-2.630	4.408
Two by Two	-8.130	4.378	-4.446	1.402
Two and Three	-3.695	0.680	-2.090	0.220

Table 1: Direction errors for the various operators

Operator	M_1	M_2	M_∞
Prewitt	0.7500	0.6882	0.5000
Sobel	0.6667	0.6569	0.5000
Roberts Cross	0.2500	0.3515	0.4970
One by Two	0.3334	0.4713	0.5006
Two by Two	0.4970	0.3514	0.2500
Two and Three	0.6981	0.6734	0.5000

Table 2: Minimum of normalized magnitude for the various operators on edge positions

4 Results

For each operator, we sampled the range of edge orientations θ, 0° to 45°, at 256 equally spaced points, and likewise sampled the range of edge offsets ρ, 0 to $\sqrt{2}/2$, at 256 equally spaced points. For each such (ρ, θ) pair we compute the magnitude and direction response, as described above. These are displayed as contour plots in Figures 4–9. (The figures for the Roberts Cross and the Two by Two operators contain the same plots, differently ordered.) In each plot, θ runs from top to bottom; ρ from left to right. The dashed curve in these plots, given by the equation $\rho = (\sin\theta + \cos\theta)/2$, delimits the region of (ρ, θ) space corresponding to edges that pass through the central unit square. Some of the contours appear jagged in these plots. This is merely an artifact produced near discontinuities by the interpolation scheme used in our contour-plotting program. The true contours are actually smooth, though they lie along a discontinuity.

A summary of these results appears in Table 1 and Table 2. For each operator, Table 1 gives the greatest negative and greatest positive value of the orientation error $\hat{\theta} - \theta$, that is, the difference between the reported edge orientation $\hat{\theta}$ (from the gradient) and the actual edge orientation θ. This table also gives the respective averages of the negative and positive values of this error $\hat{\theta} - \theta$. In Table 2, the magnitude responses of the operators using various norms are normalized by division by the greatest magnitude achieved, to make these magnitudes comparable across the various operators. The table shows the normalized minimum values attained by the magnitude on edge positions. It gives bounds on the variation in reported magnitude caused by the positioning of an edge with respect to the pixel grid.

All results presented in the figures and tables so far are for edge positions only. However, the

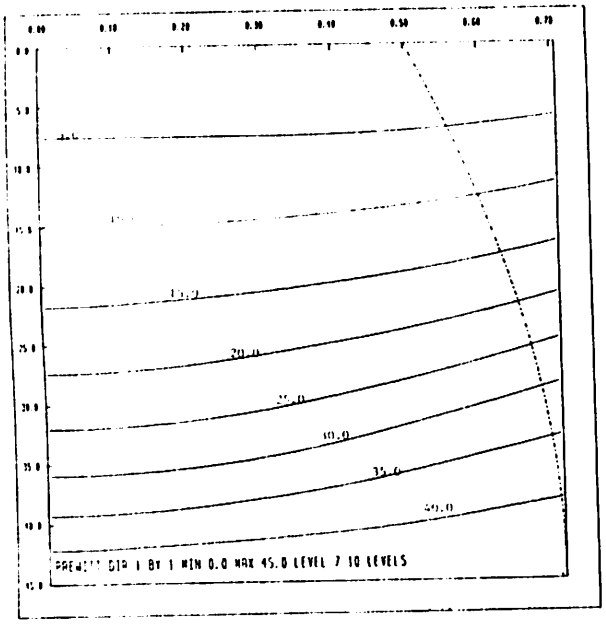

(a)

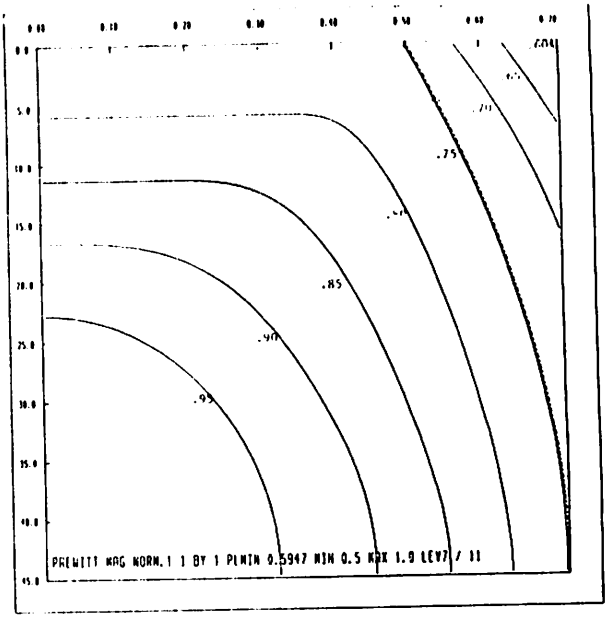

(b)

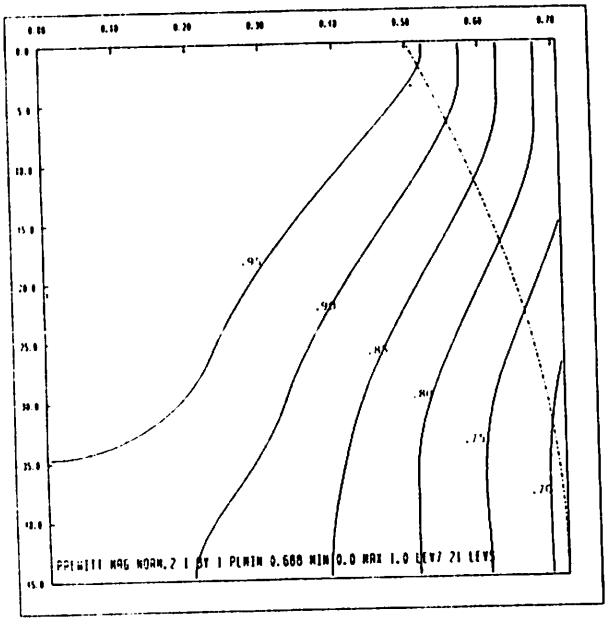

(c)

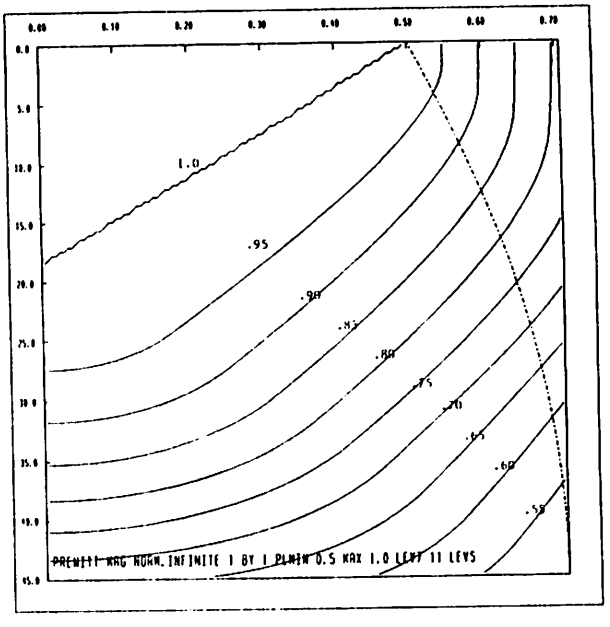

(d)

Figure 4: Prewitt operator response: (a.) reported edge orientation $\hat{\theta}$ (b.) edge magnitude M_1 (c.) edge magnitude M_2 (d.) edge magnitude M_∞ (a. to d. read ρ from left to right, θ from top to bottom)

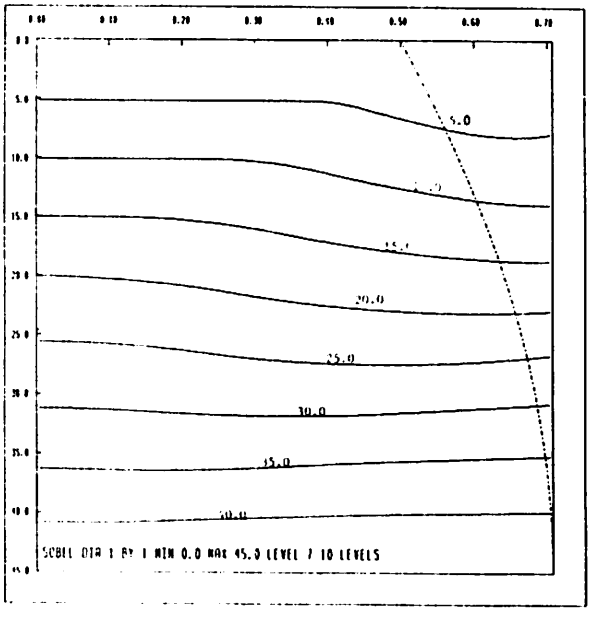

(a)

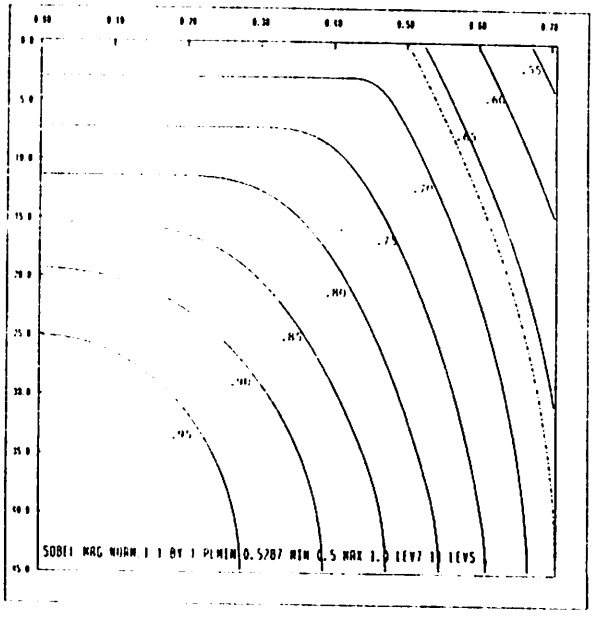

(b)

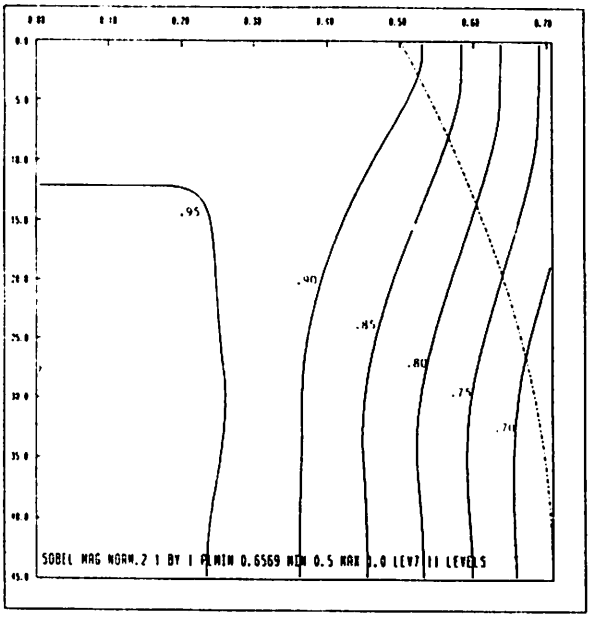

(c)

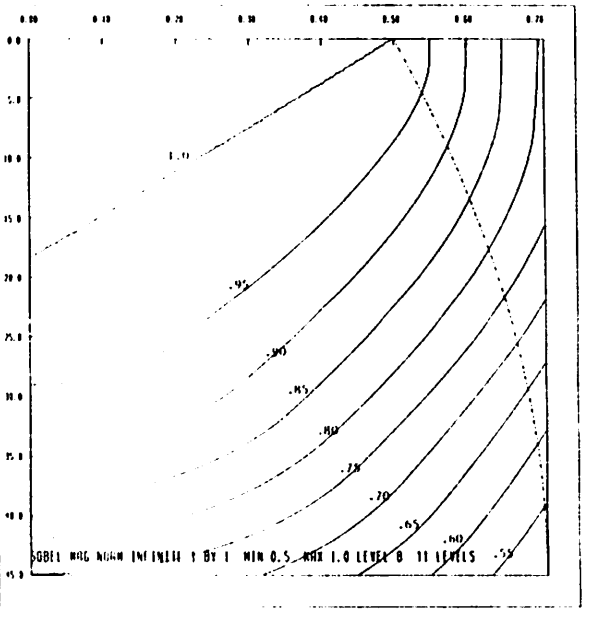

(d)

Figure 5: Sobel operator response: (a.) edge orientation $\hat{\theta}$ (b.) edge magnitude M_1 (c.) edge magnitude M_2 (d.) edge magnitude M_∞ (a. to d. read ρ from left to right, θ from top to bottom)

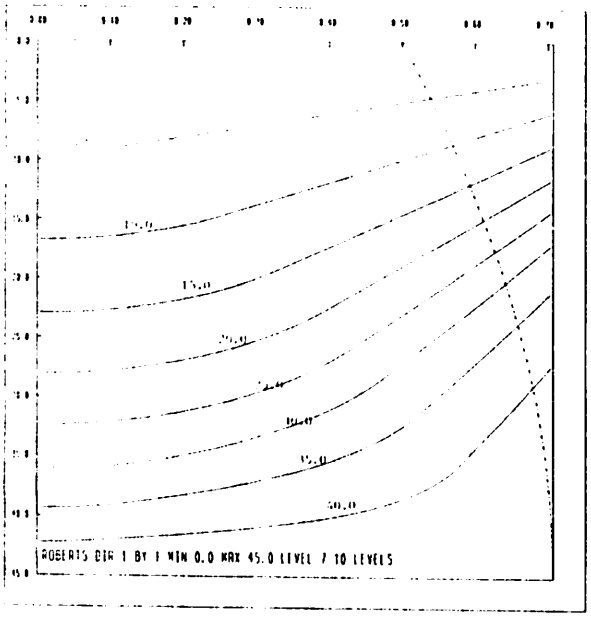

(a)

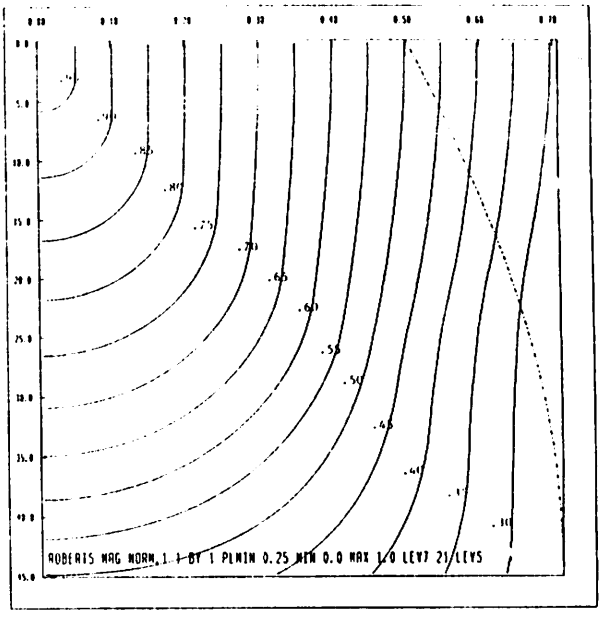

(b)

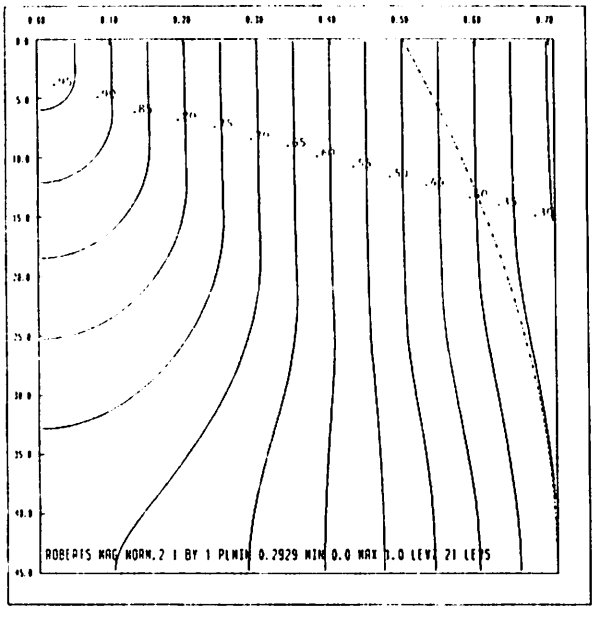

(c)

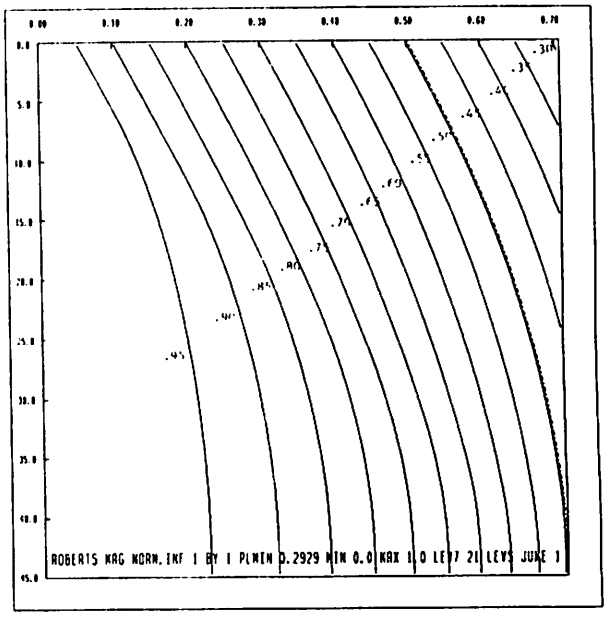

(d)

Figure 6: Roberts Cross operator response: (a.) edge orientation $\hat{\theta}$ (b.) edge magnitude M_1 (c.) edge magnitude M_2 (d.) edge magnitude M_∞ (a. to d. read ρ from left to right, θ from top to bottom)

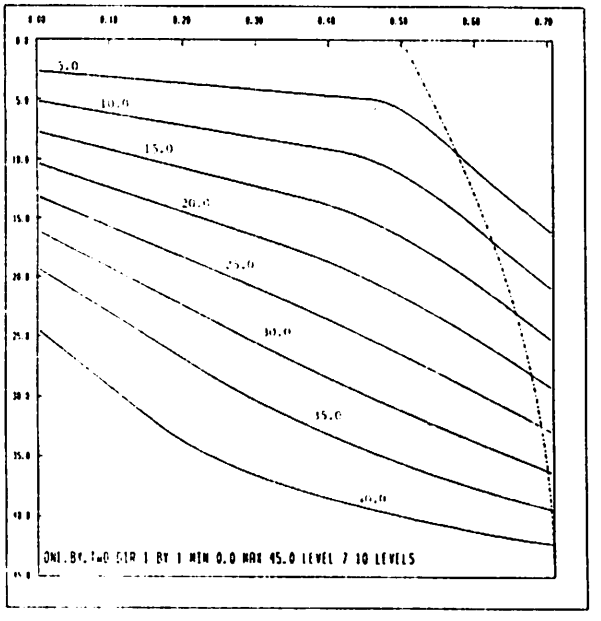

(a)

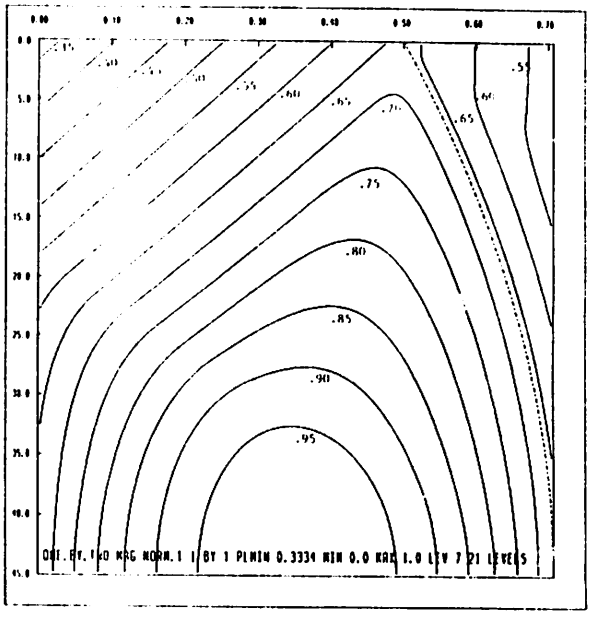

(b)

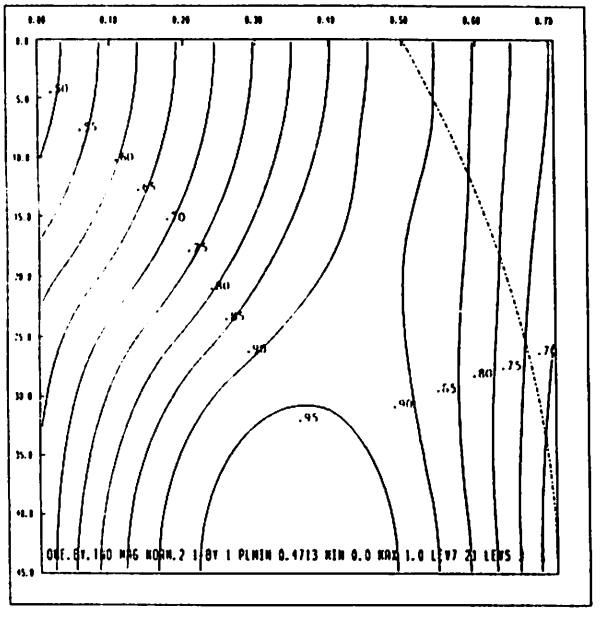

(c)

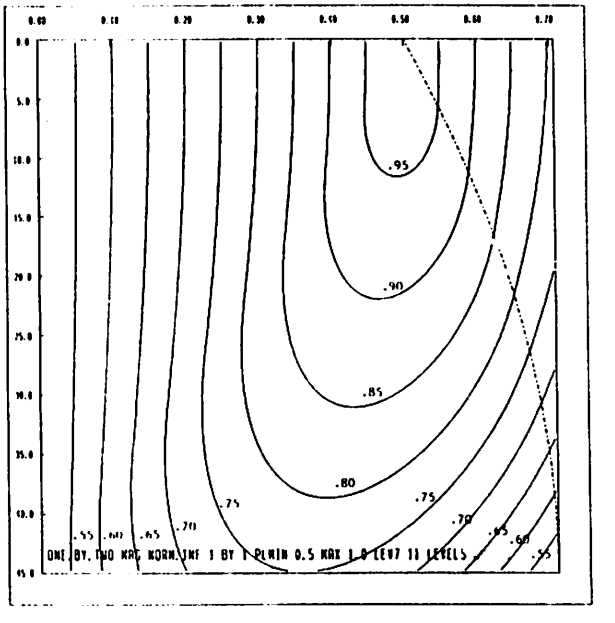

(d)

Figure 7: "One by Two" operator response: (a.) edge orientation $\hat{\theta}$ (b.) edge magnitude M_1 (c.) edge magnitude M_2 (d.) edge magnitude M_∞ (a. to d. read ρ from left to right, θ from top to bottom)

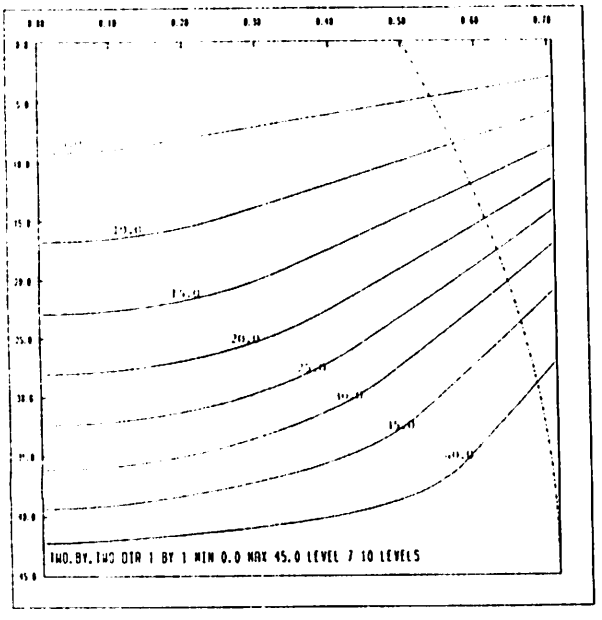

(a)

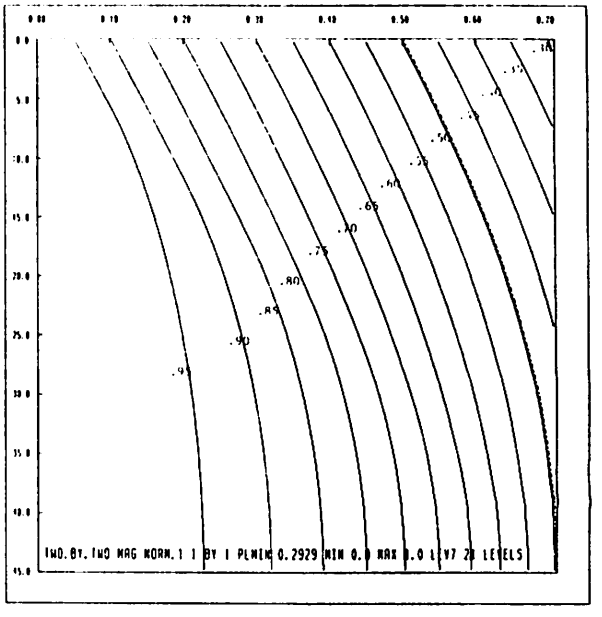

(b)

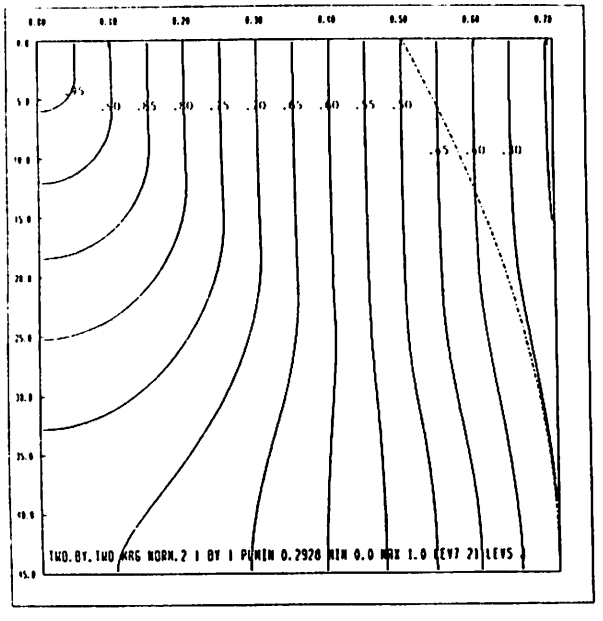

(c)

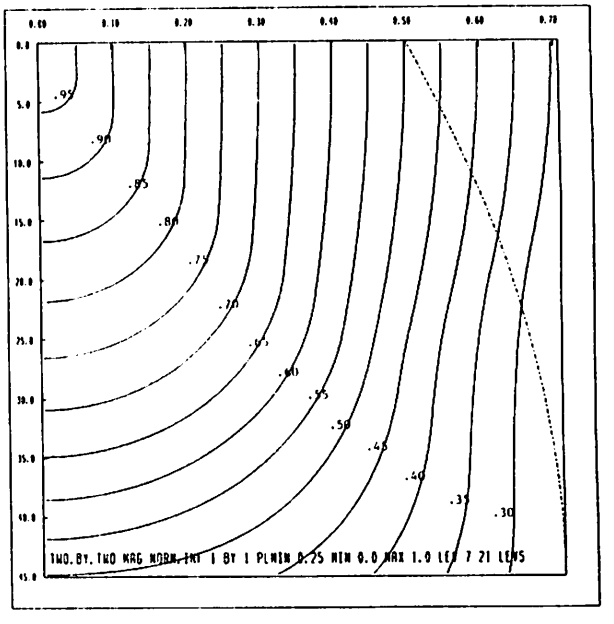

(d)

Figure 8: "Two by Two" operator response: (a.) edge orientation $\hat{\theta}$ (b.) edge magnitude M_1 (c.) edge magnitude M_2 (d.) edge magnitude M_∞ (a. to d. read ρ from left to right, θ from top to bottom)

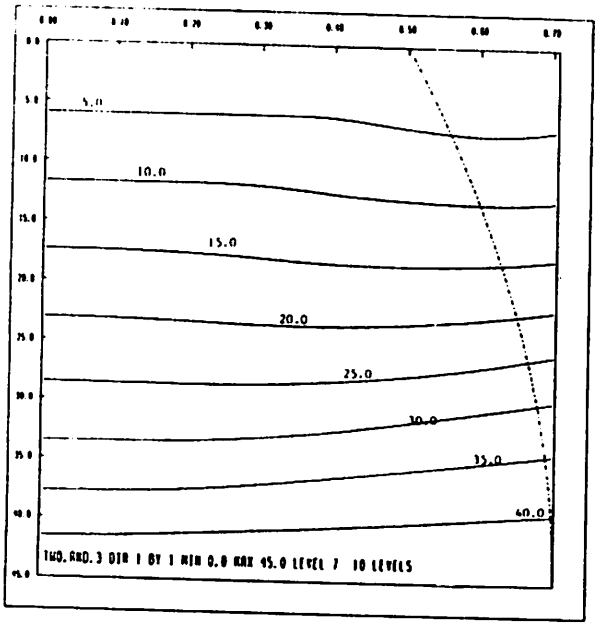

(a)

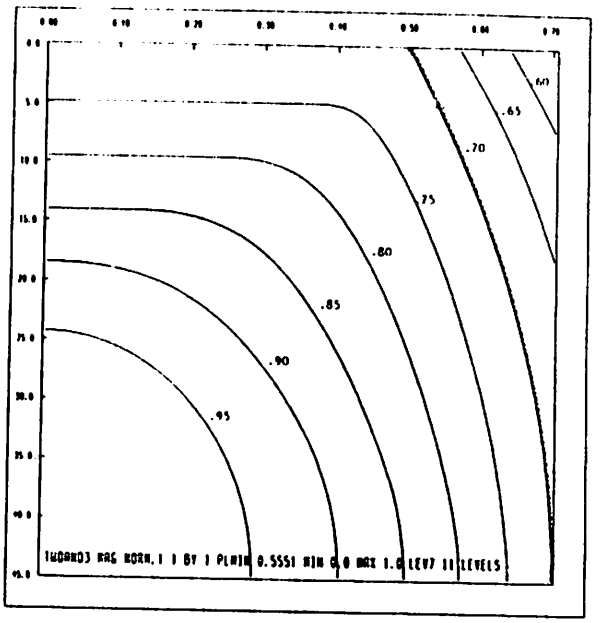

(b)

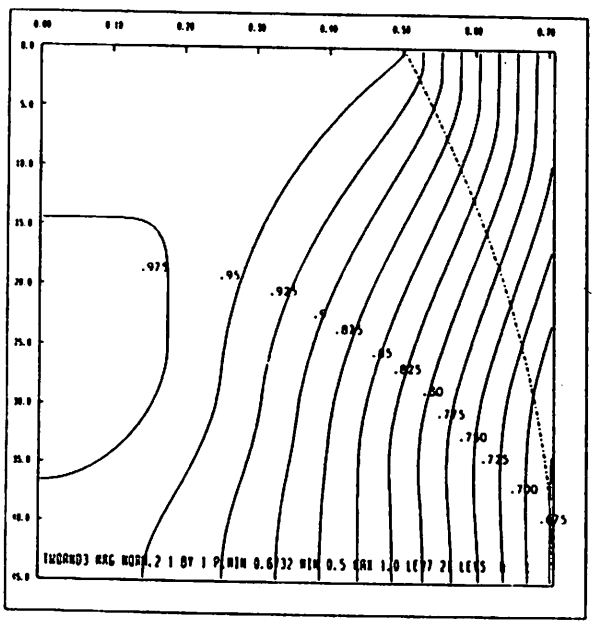

(c)

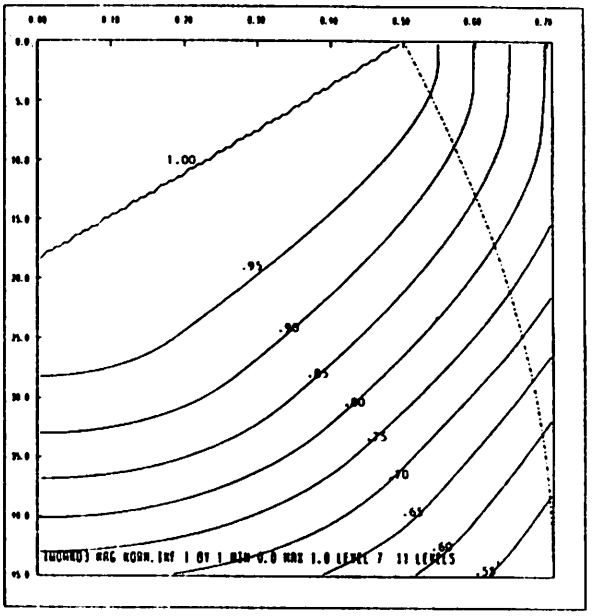

(d)

Figure 9: "Two and Three" operator response: (a.) edge orientation $\hat{\theta}$ (b.) edge magnitude M_1 (c.) edge magnitude M_2 (d.) edge magnitude M_∞ (a. to d. read ρ from left to right, θ from top to bottom)

entire response of an operator to any edge passing through its $n \times n$ neighborhood is also of interest. Figure 10 shows this extended response, in direction and magnitude, for the Sobel operator, with θ again sampled at 256 equally spaced points in the range 0° to 45°, but with ρ sampled at 256 equally spaced points in the range 0 to $n\sqrt{2}/2$, that is $3\sqrt{2}/2$. The two dashed curves in this figure delimit edges passing respectively through the central unit square and through the whole 3×3 neighborhood. Outside the central unit square, the responses of the other operators tend to be much the same as that of the Sobel operator, and are therefore not presented here.

5 Discussion

There are a number of observations that can be made from the results presented here. One immediate observation is that the Roberts Cross and Two by Two operators are in a sense duals of each other; their responses are identical, except with M_1 and M_∞ interchanged. This was unexpected, but in retrospect, hardly surprising. Both take simple differences between adjacent pixels, using axis directions tilted at an angle of 45° from each other. However, the interaction between the axis tilt and the pixel spacing and digitization is not trivial, so this duality is an interesting result.

Other observations having to do with reported direction and reported magnitude are presented below. Since the observations about magnitude are somewhat simpler, we will deal with them first.

5.1 Edge magnitudes

It is clear from Figures 4–10 and Table 2 that even for a straight step edge of uniform constrast, the magnitude reported by an edge operator at edge positions along that edge can vary considerably. This poses difficulties for any scheme which attempts to distinguish edge positions from non-edge positions by using a threshold on magnitude.

Consider the magnitude plots in Figure 5. The dashed boundary curve separates responses on edge positions from responses on non-edge positions. Since the magnitude contours cut *across* this curve, it is clearly impossible to choose a magnitude threshold that will discriminate edge positions from non-edge positions for edges at arbitrary orientation, even if the edges are of uniform contrast. Any threshold will include non-edge positions for some orientations and exclude edge positions for other orientations. Only for the Roberts Cross with the L_∞ norm, and the Two by Two with the L_1 norm (which is essentially identical) do the magnitude contours run parallel to the boundary curve, and only for the Prewitt, Sobel, and "Two-and-three" with the L_1 norm do the magnitude contours not cross the boundary curve (though these latter operators have a strong orientation bias in their responses). Only for these edge operators is it possible to choose a threshold that will discriminate edge positions from non-edge positions for edges of uniform contrast with arbitrary orientation. Notice that this "preferred" behavior occurs for the L_1 and L_∞ norms, which are often regarded merely as computationally expedient approximations to the L_2 norm. Here we have situations in which the simpler norm has an advantage that is not just computational.

It is often felt to be desirable that the response of an edge operator should show no orientation bias, that is, the magnitude reported by the operator should be independent of edge orientation. Diagrammatically, this means that the contour lines should run vertically in the figures. Notice that for the 3×3 operators there are regions where the contour lines run *horizontally*, meaning that the magnitude depends *solely* on the orientation for an edge of fixed contrast . Also, the requirement that an edge operator have no orientation bias and the requirement that it be possible

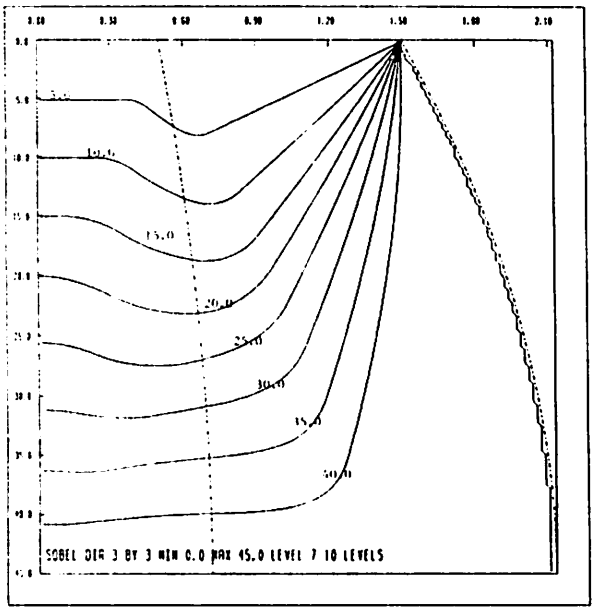

(a)

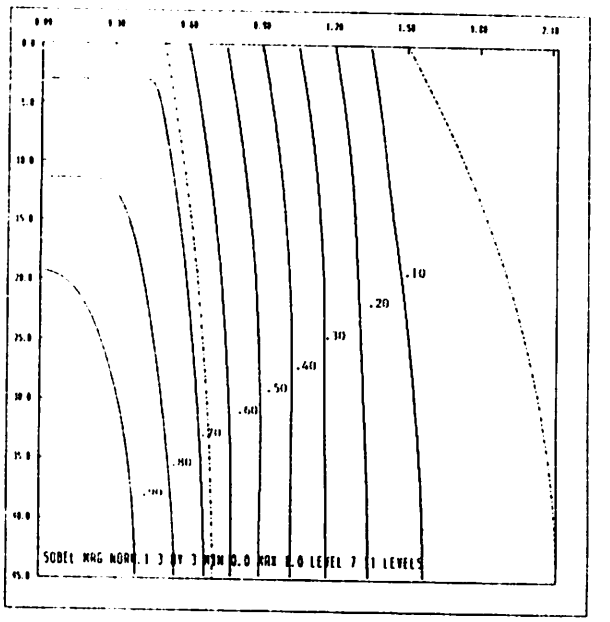

(b)

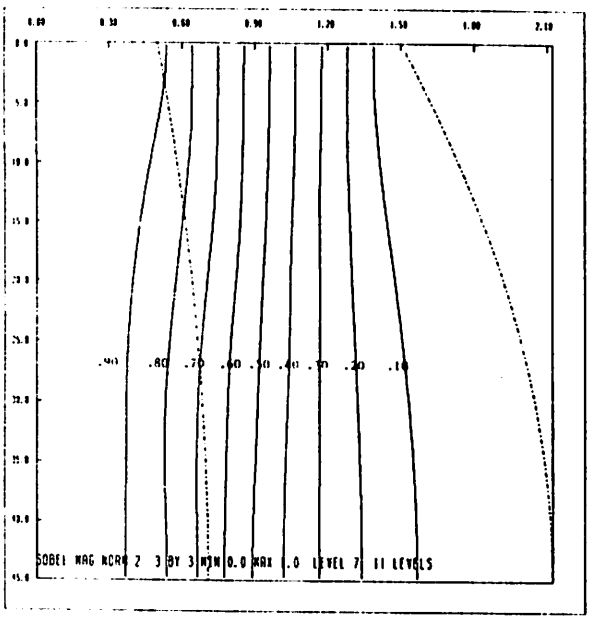

(c)

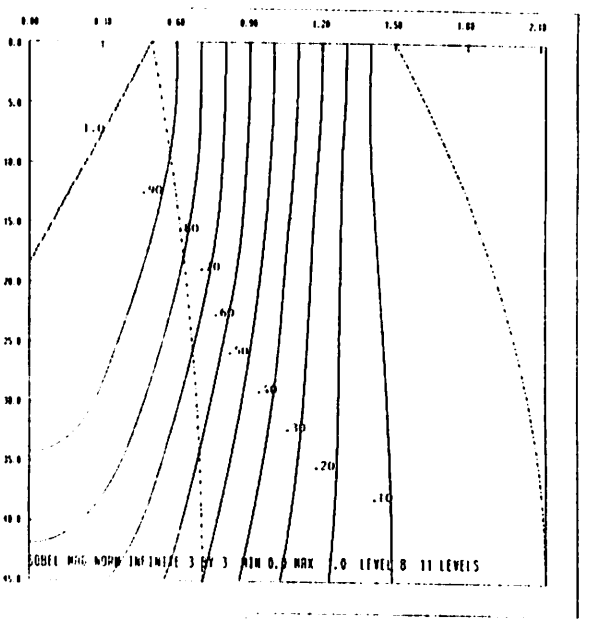

(d)

Figure 10: Sobel operator response on full 3×3 neighborhood: (a.) edge orientation $\hat{\theta}$ (b.) edge magnitude M_1 (c.) edge magnitude M_2 (d.) edge magnitude M_∞ (a. to d. read ρ from left to right, θ from top to bottom)

to threshold edge positions from non-edge positions are incompatible, in that the former requires that the contours be everywhere vertical while the latter requires that locally they run parallel to the (dashed) boundary curve, which is not vertical.

We should point out that this discussion ignores consideration of noise. The problem of separating edge positions from non-edge positions associated with the same edge, and from false edges caused by noise is a complex one, needing further study. One standard approach is to use not just isolated magnitude responses, but instead some of the structural properties of an edge, like relative magnitude of directionally adjacent pixels, as is done in [13]. However, our treatment does isolate the difficulties caused purely by the geometry of the edge operators, and shows that they are far from negligible. It can also seen that when noise is small, these geometric difficulties with thresholding can be overcome by suitable choice of operator and norm.

5.2 Edge orientation

We now discuss the orientation response of the operators. Under the conditions we have considered, the absolute error in direction reported by these operators at edge positions can be quite large, greater than 8° for the Roberts Cross and Two by Two operators. Even the best of the commonly used operators, the Sobel, has a directional error of almost 4°.

Suspecting that the superior performance of the Sobel operator was due to its heavier weighting of pixels closer to the center of the neighborhood, we experimented with a number of other combinations of mask weights. The Two and Three operator was the simplest and best of these. Even so, the improvement over the Sobel in worst-case error was very small, and its average error was worse, suggesting that little is to be gained by further adjusting weights.

Consideration of the contour plots in Figures 4–9 shows that the raw orientation errors of Table 1 are not so important of themselves. If a given operator reports an edge orientation $\hat{\theta}$, all we can say is that the true values of θ and ρ lie along the contour defined by $\hat{\theta}$, but they are otherwise undetermined. However, if $\hat{\theta}$ depends monotonically on θ alone (this would be manifest by perfectly horizontal contours), then this dependence could be inverted to yield the true orientation θ from the reported orientation $\hat{\theta}$. For a number of the operators, the contours are not quite horizontal, suggesting an approximate calibration procedure instead. For each reported orientation $\hat{\theta}$ there is a small range of possible true orientations θ that could have yielded $\hat{\theta}$. Choose, say, the midpoint of that range as an estimate of θ. We have not yet implemented this idea, but it seems that for most of the operators it could considerably reduce the orientation error, perhaps to as little as 1°. In practice, the mapping could be precomputed, and stored in a lookup table, making the calibration process very inexpensive.

An even more intriguing idea is to make use of the orientation outputs from two different operators at the same position. If one operator reports orientation $\hat{\theta}_1$, then we know that the true ρ and θ must lie along the $\hat{\theta}_1$ contour of that operator. Similarly, if a second operator reports $\hat{\theta}_2$, we know that ρ and θ must lie on the $\hat{\theta}_2$ contour of that second operator. If these two contour lines have a unique intersection, we can read off the true orientation θ, and the true sub-pixel edge offset ρ, and from these and the magnitude determine the actual step height. Thus the responses of two different operators yield all the parameters of the edge, including its location to sub-pixel precision.

Again, the contour intersections could be precomputed, and stored in a two-dimensional lookup table, indexed by $\hat{\theta}_1$ and $\hat{\theta}_2$. Further savings could be realised by doing these computations in terms

of the ratios I_y/I_x for the two operators, avoiding the inverse tangent. Even though it is necessary to apply two different operators, this may still be cheaper than applying a single expensive operator. Preliminary examination shows that several pairs of operators might be suitable for this technique, in that their sets of orientation contours have well defined intersections.

It should be kept in mind that these results are for an ideal case of response to a perfect, noise-free, straight step edge, in which it is assumed that edge positions can be distinguished from non-edge positions. In practice, these directional errors will be much larger, and may make infeasible the use of the calibration or contour-intersection techniques described above. Even for perfect edges, the directional error for non-edge positions can be as great as 45°, as can be seen in Figure 10. This is because the orientation reported for an edge that just clips the neighborhood will be 45°, regardless of the edge's actual orientation—even if its orientation is close to 0°.

Clearly, all these observations are crucial to techniques such as [4,19] which use gradient directions for straight-edge grouping and other purposes. On the one hand they set limits to the performance that can be expected from these techniques using these conventional simple edge operators; on the other hand they suggest computationally inexpensive ways of markedly improving such techniques.

6 Future Work

There are four directions in which this work could be developed. First is to examine other edge and digitization models. A suitable digitization model would use Gaussian weighted receptive fields for pixels. This could be somewhat easier to analyse than the model with square receptive fields used here, and would also provide a more realistic approximation to the optical characteristics of real cameras. Other edge models could be considered. However, even curved edges should appear almost straight at the scale of the neighborhoods considered; ramp edges could be modelled by blurring of a step edge, which could be absorbed into the Gaussian weighting of the pixel receptive fields. Therefore it seems unlikely that any more complex edge model would be necessary. As mentioned earlier, we chose step edges and square receptive fields in this paper chiefly for consistency with previous work [1,2].

A second direction is to consider the effects of image intensity quantization and noise. It is clear from our analysis here, and from our experience with actual images, that the effects of spatial discretization alone are quite significant, and certainly comparable to those caused by noise and intensity quantization. However it would be desirable to quantify the relative importance of these effects more rigorously.

A third direction is to carry out some of this analysis in the spatial-frequency domain. Except for the computation of the magnitude and direction of the gradient from its I_x and I_y components, all steps of processing are amenable to such an analysis, which would perhaps lead to further insights into the behavior of edge operators under the conditions we have considered.

And the fourth direction is to implement and evaluate the techniques we have proposed (calibration and contour-intersection) for improving the performance of edge operators. While our preliminary examination indicates that such techniques will be feasible, it remains to prove them in practice, and to quantify the improvements obtainable.

7 Conclusion

We have mapped out the magnitude and orientation responses of a number of commonly used simple edge operators on a step edge, using a digitization model in which pixels have uniform square receptive fields. This allows us to quantify the limits to the performance of these operators, even under ideal conditions. For computing edge orientations, it would seem that the Sobel operator is most accurate; for discriminating edges from non-edges by thresholding, only the Roberts Cross operator (using the L_∞ norm), the Two by Two operator (using the L_1 norm), or possibly one of the 3×3 operators (using the L_1 norm) would be suitable. This work also suggests several computationally simple techniques (calibration and contour intersection) for improving the performance of edge operators.

8 Acknowledgements

Thanks to George Reynolds, Richard Weiss, and Al Hanson for reading drafts, and to Laurie Waskiewicz for help with preparing figures. The contour-intersection technique arose from an observation by Lance Williams.

References

[1] I. E. Abdou, "Quantitative methods of edge detection", Image Processing Institute, University of Southern California, Los Angeles, CA, USCIPI Report 630, 1978.

[2] I. E. Abdou and W. K. Pratt, "Quantitative design and evaluation of enhancement/thresholding edge detectors", *Proc. IEEE*, vol. 67, no. 5, pp. 753-763, May 1979.

[3] D. J. Bryant and D. W. Bouldin, "Evaluation of edge operators using relative and absolute grading", *Proc. IEEE Conf. Pattern Recognition and Image Processing*, Chicago, IL, pp. 138-145, Aug. 1979.

[4] J. B. Burns and A. R. Hanson and E. M. Hanson, "Extracting Straight Lines", *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 8, no. 4, pp 425-456, July 1986.

[5] J. F. Canny, "A Computational Approach to Edge Detection", *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. PAMI-8, November 1986, pp. 679-698.

[6] L. S. Davis, "A Survey of Edge Detection Techniques", *Computer Graphics and Image Processing*, vol. 4, no. 3, pp. 248-270, September 1975.

[7] R. O. Duda and P. E. Hart, *Pattern Classification and Scene Analysis*, New York: Wiley, p. 271, 1973.

[8] J. R. Fram and E. S. Deutsch, "A quantitative study of the orientation bias of some edge detector schemes", University of Maryland, College Park MD, Computer Science Technical Report 285, Jan. 1974.

of the ratios I_y/I_x for the two operators, avoiding the inverse tangent. Even though it is necessary to apply two different operators, this may still be cheaper than applying a single expensive operator. Preliminary examination shows that several pairs of operators might be suitable for this technique, in that their sets of orientation contours have well defined intersections.

It should be kept in mind that these results are for an ideal case of response to a perfect, noise-free, straight step edge, in which it is assumed that edge positions can be distinguished from non-edge positions. In practice, these directional errors will be much larger, and may make infeasible the use of the calibration or contour-intersection techniques described above. Even for perfect edges, the directional error for non-edge positions can be as great as 45°, as can be seen in Figure 10. This is because the orientation reported for an edge that just clips the neighborhood will be 45°, regardless of the edge's actual orientation—even if its orientation is close to 0°.

Clearly, all these observations are crucial to techniques such as [4,19] which use gradient directions for straight-edge grouping and other purposes. On the one hand they set limits to the performance that can be expected from these techniques using these conventional simple edge operators; on the other hand they suggest computationally inexpensive ways of markedly improving such techniques.

6 Future Work

There are four directions in which this work could be developed. First is to examine other edge and digitization models. A suitable digitization model would use Gaussian weighted receptive fields for pixels. This could be somewhat easier to analyse than the model with square receptive fields used here, and would also provide a more realistic approximation to the optical characteristics of real cameras. Other edge models could be considered. However, even curved edges should appear almost straight at the scale of the neighborhoods considered; ramp edges could be modelled by blurring of a step edge, which could be absorbed into the Gaussian weighting of the pixel receptive fields. Therefore it seems unlikely that any more complex edge model would be necessary. As mentioned earlier, we chose step edges and square receptive fields in this paper chiefly for consistency with previous work [1,2].

A second direction is to consider the effects of image intensity quantization and noise. It is clear from our analysis here, and from our experience with actual images, that the effects of spatial discretization alone are quite significant, and certainly comparable to those caused by noise and intensity quantization. However it would be desirable to quantify the relative importance of these effects more rigorously.

A third direction is to carry out some of this analysis in the spatial-frequency domain. Except for the computation of the magnitude and direction of the gradient from its I_x and I_y components, all steps of processing are amenable to such an analysis, which would perhaps lead to further insights into the behavior of edge operators under the conditions we have considered.

And the fourth direction is to implement and evaluate the techniques we have proposed (calibration and contour-intersection) for improving the performance of edge operators. While our preliminary examination indicates that such techniques will be feasible, it remains to prove them in practice, and to quantify the improvements obtainable.

7 Conclusion

We have mapped out the magnitude and orientation responses of a number of commonly used simple edge operators on a step edge, using a digitization model in which pixels have uniform square receptive fields. This allows us to quantify the limits to the performance of these operators, even under ideal conditions. For computing edge orientations, it would seem that the Sobel operator is most accurate; for discriminating edges from non-edges by thresholding, only the Roberts Cross operator (using the L_∞ norm), the Two by Two operator (using the L_1 norm), or possibly one of the 3×3 operators (using the L_1 norm) would be suitable. This work also suggests several computationally simple techniques (calibration and contour intersection) for improving the performance of edge operators.

8 Acknowledgements

Thanks to George Reynolds, Richard Weiss, and Al Hanson for reading drafts, and to Laurie Waskiewicz for help with preparing figures. The contour-intersection technique arose from an observation by Lance Williams.

References

[1] I. E. Abdou, "Quantitative methods of edge detection", Image Processing Institute, University of Southern California, Los Angeles, CA, USCIPI Report 630, 1978.

[2] I. E. Abdou and W. K. Pratt, "Quantitative design and evaluation of enhancement/thresholding edge detectors", *Proc. IEEE*, vol. 67, no. 5, pp. 753–763, May 1979.

[3] D. J. Bryant and D. W. Bouldin, "Evaluation of edge operators using relative and absolute grading", *Proc. IEEE Conf. Pattern Recognition and Image Processing*, Chicago, IL, pp. 138–145, Aug. 1979.

[4] J. B. Burns and A. R. Hanson and E. M. Hanson, "Extracting Straight Lines", *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 8, no. 4, pp 425–456, July 1986.

[5] J. F. Canny, "A Computational Approach to Edge Detection", *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. PAMI–8, November 1986, pp. 679–698.

[6] L. S. Davis, "A Survey of Edge Detection Techniques", *Computer Graphics and Image Processing*, vol. 4, no. 3, pp. 248–270, September 1975.

[7] R. O. Duda and P. E. Hart, *Pattern Classification and Scene Analysis*, New York: Wiley, p. 271, 1973.

[8] J. R. Fram and E. S. Deutsch, "A quantitative study of the orientation bias of some edge detector schemes", University of Maryland, College Park MD, Computer Science Technical Report 285, Jan. 1974.

[9] J. R. Fram and E. S. Deutsch, "On the quantitative evaluation of edge detection schemes and their comparison with human performance", *IEEE Trans. Computing*, vol. C-24, no. 6, June 1975, pp. 616–628.

[10] R. M. Haralick, "Digital Step Edges from Zero Crossings of Second Directional Derivatives", *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. PAMI-6, January 1984, pp. 58–68.

[11] M. Hueckel, "An operator which locates edges in digitized pictures", *J. ACM*, vol. 18, no. 1, January 1971, pp. 113–125.

[12] M. Hueckel, "A local visual operator which recognizes edges and lines", *J. ACM*, vol. 20, no. 4, October 1973, pp. 634–647.

[13] L. J. Kitchen and A. Rosenfeld, "Non-maximum suppression of edge gradients makes them easier to threshold", *Pattern Recognition Letters*, vol. 1, no. 2, December 1982, pp. 93–94.

[14] L. J. Kitchen and A. Rosenfeld, "Edge evaluation using local edge coherence", *IEEE Trans. Systems, Man, and Cybernetics*, vol. SMC-11, no. 9, September 1981, pp. 597–605.

[15] D. Marr and E. Hildreth, "Theory of edge detection", *Proc. Roy. Soc. London*, Ser. B, vol. 207, 1980, pp. 187–217.

[16] L. Mero and Z. Vassy, "A simplified and fast version of the Hueckel operator for finding optimal edges in pictures", *Proc., 4th IJCAI*, September 1975, pp. 650–655.

[17] V. S. Nalwa and T. O. Binford, "On Detecting Edges", *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. PAMI-8, November 1986, pp. 679–698.

[18] R. Nevatia, "Evaluation of a simplified Hueckel edge-line detector", Note, *CGIP*, December 1977, pp. 582–588.

[19] S. Peleg, "Straight edge enhancement and mapping", University of Maryland, College Park MD, Computer Science Technical Report 694, Sept. 1978.

[20] W. K. Pratt, *Digital Image Processing*, Wiley, New York, 1978, pp. 495–501.

[21] J. M. S. Prewitt, "Object enhancement and extraction", in, *Picture Processing and Psychopictorics*, B.S. Lipkin and A. Rosenfeld (Eds.), New York: Academic Press, 1970.

[22] L. G. Roberts, "Machine perception of three-dimensional solids", in *Optical and Electro-optical Information Processing*, J. P. Tippett et al. (Eds.), Cambridge, MA: MIT Press, 1965.

[23] A. Rosenfeld and A. C. Kak, *Digital Picture Processing*, New York: Academic Press, 1976.