[29] H. K. Nishihara and N. G. Larson, "Towards a real time implementation of the Marr and Poggio stereo matcher," in *Proc. ARPA Image Understanding Workshop*, 1981, pp. 114–120.
[30] J. Prewitt, "Object enhancement and extraction," in *Picture Processing and Psychopictorics*, B. Lipkin and A. Rosenfeld, Eds. New York: Academic, 1970, pp. 75–149.
[31] V. Torre and T. Poggio, "On edge detection," Artif. Intell. Lab., Mass. Inst. Technol., Memo. 768, Oct. 1983.
[32] A. P. Witkin, "Recovering surface shape and orientation from texture," *Artif. Intell.*, vol. 17, pp. 17–45, 1981.
[33] A. L. Yuille and T. Poggio, "Scaling theorems for zero-crossings," Artif. Intell. Lab., Mass. Inst. Technol., Memo. 772, June 1983.
[34] —, "Fingerprints theorems for zero-crossings," Artif. Intell. Lab., Mass. Inst. Technol., Memo. 730, Oct. 1983.
[35] S. W. Zucker, "Motion and the Mueller-Lyer illusion," Dep. Elec. Eng., McGill Univ., Montreal, P.Q., Canada, Tech. Rep. 80-2R, 1980.
[36] —, "Computer vision and human perception: An essay on the discovery of constraints," in *Proc. 7th Int. Joint Conf. Artif. Intell.*, Vancouver, B.C., Canada, vol. 2, Aug. 1981, pp. 1102–1116.

Author's Reply[2]

R. M. HARALICK

***Abstract*—We present evidence that the Laplacian zero-crossing operator does not use neighborhood information as effectively as the second directional derivative edge operator. We show that the use of a Gaussian smoother with standard deviation 5.0 for the Laplacian of a Gaussian edge operator with a neighborhood size of 50 × 50 both misses and misplaces edges on an aerial image of a mobile home park. Contrary to Grimson and Hildreth's results, our results of the Laplacian edge detector on a noisy test checkerboard image are also not as good as the second directional derivative edge operator. We conclude by discussing a number of open issues on edge operator evaluation.**

I. Introduction

Grimson and Hildreth [3] suggest that comparisons between edge detectors should be done without regard to considerations of neighborhood size. Their suggestion for an edge detector is to eliminate noise on the input image by smoothing with a sufficiently broad Gaussian filter, take the Laplacian of the smoothed image, and mark pixels as edges if in some direction the pixel on the convolved image has a zero-crossing with a high enough slope. They state that for the test checkerboard image with 20 × 20 checks and a check contrast-to-noise ratio of 2 : 1 using a Gaussian smoother with standard deviation of 5.0, the probability of a true edge being assigned an edge by their edge detector is about 0.9 when the zero crossing slope is given a threshold in a way which equalizes the number of true edges assigned as nonedges with the number of nonedges assigned as edges. They argue for a neighborhood size in the range of 45 × 45 rather than the truncated neighborhood size of 11 × 11 used in Haralick.[1]

Although Grimson and Hildreth [3] do not mention it in their correspondence, they did, in private correspondence, note that the equation given by Haralick[1] for the Laplacian of a Gaussian consistently had a typographical error of a misplaced parenthesis. Computer programs and results, however, were correct.

We attempted to replicate the Grimson–Hildreth result using a Gaussian smoother with standard deviation of 5.0 with a neighborhood size of 50 × 50. Any pixel which had a zero crossing slope greater than 10 zero-crossings of the smoothed Laplacian was assigned an edge. True edges were declared for any pixel of the no-noise checkerboard, which was black but bordered a white pixel, or which was a white pixel and bordered a black pixel. Our results indicate that, given a pixel is a true edge, the probability that the pixel is assigned an edge is 0.7217. Given that a pixel is assigned as an edge, the probability that it is a true edge pixel is 0.7155. This differs considerably from their result. It would be worthwhile to carefully review each of our procedures to determine why this difference arises. Is it due to a different definition of true edge? Is it due to a difference in the zero-crossing slope computation?

Even if the replication agreed with the Grimson–Hildreth experiment, the situation would be more complicated than it appears on the surface. From a signal content/noise content point of view, the standard deviation of the Gaussian filter must be set based on the size distribution of the homogeneous regions, their relative contrasts, and the amount of noise. A standard deviation of 5.0 for a Gaussian averager may leave objects such as the 20 × 20 checks intact, but would tend to smooth out of existence objects which are small or thin. Thus, there are circumstances in which a standard deviation of 5.0 would be inappropriately large, and it is precisely for this reason that a fixed window size was selected to do the experiments.

To see the folly of not fixing the size of the window, consider an image whose size is as large as we like, whose left-hand side is noisy black, and whose right-hand side is noisy white. Suppose the signal-to-noise ratio is reasonable. Under these circumstances, consider how we would want to evaluate edge operators. Since the geometry is utterly simple and the objects are as large as we would like, each edge operator proponent could find a window of sufficiently large size so that the edge operator produces a result of prespecified accuracy. Obviously, in this situation the above evaluation is meaningless. What we must do is perform the evaluation under conditions in which the pixel information provided to the edge operator is limited and then perform the evaluation under the limiting information conditions. Under these circumstances, an edge operator could be said to be uniformly better than other edge operators if under each possible information limiting condition it performs better than all the other edge operators. Thus, performance in controlled experiments must be performance as a function of information utilized. The key issue is, how well does the operator utilize a fixed information set?

II. Experiments

To show the problem of an excessively large standard deviation for the Gaussian smoother, we try to determine the edges of the aerial image of a mobile home park, shown in Fig. 1. We perform three experiments. In the first experiment, a Gaussian standard deviation of 5.0 is used with an adequate 45 × 45 window as the smoother preceding the Laplacian. The zero-crossings obtained having a nonzero slope are shown in Fig. 2. Notice how many edges are not detected and that many edges are misplaced around nearly straight boundaries as well as around corners. This is only a reasonable edge image if the rows of the mobile homes are the desired objects. It is not a reasonable edge image if the boundaries of the individual homes are desired.

In the second experiment, a Gaussian standard deviation of 0.8 is used with an adequate 7 × 7 window as the smoother preceding the Laplacian. The zero-crossings obtained having a slope greater than 2 are shown in Fig. 3. 25 percent of the pixels are assigned edges. Although noisy, at least this image shows the individual edges around the mobile homes.

The third experiment uses the second directional derivative zero-crossing edge operator. The equally weighted least squares bivariate cubic fit is done in a 7 × 7 neighborhood, and a pixel

[2]Manuscript received March 27, 1984; revised September 10, 1984.
The author is with Machine Vision International, Ann Arbor, MI 48104.

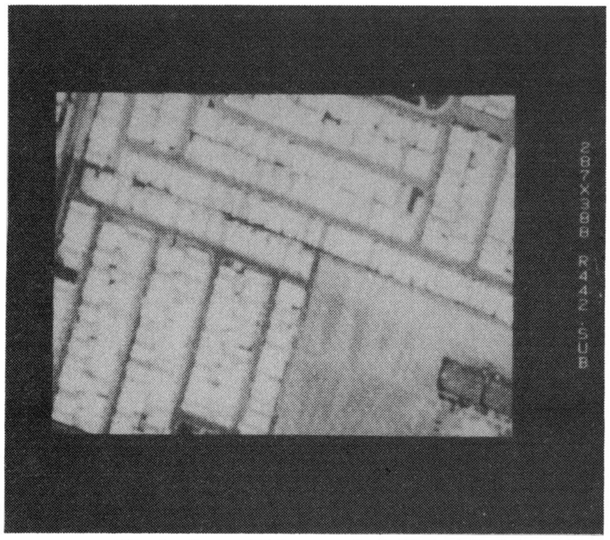

Fig. 1. An aerial image of a trailer park.

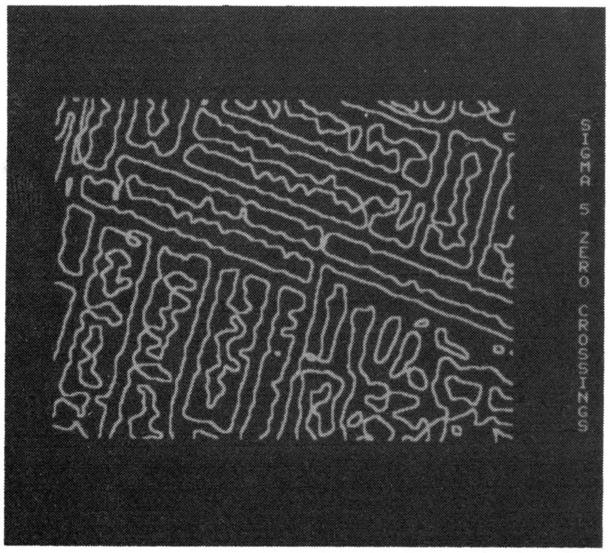

Fig. 2. The zero-crossings of a Laplacian edge detector having a Gaussian standard deviation of 5.0 and using a window of 45 X 45. 22 percent of the pixels are assigned as edges.

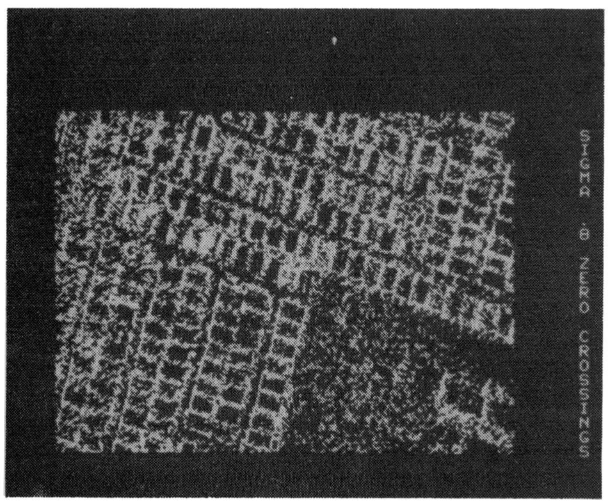

Fig. 3. The zero-crossings of a Laplacian edge detector having a Gaussian standard deviation of 0.8 and using a window of 7 X 7. 25 percent of the pixels are assigned as edges.

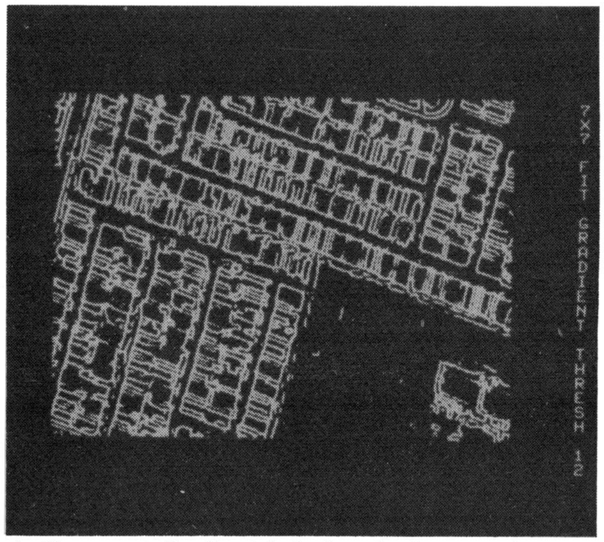

Fig. 4. The second directional derivative edge detector using an equally weighted cubic fit in a 7 X 7 window. 25 percent of the pixels are assigned as edges.

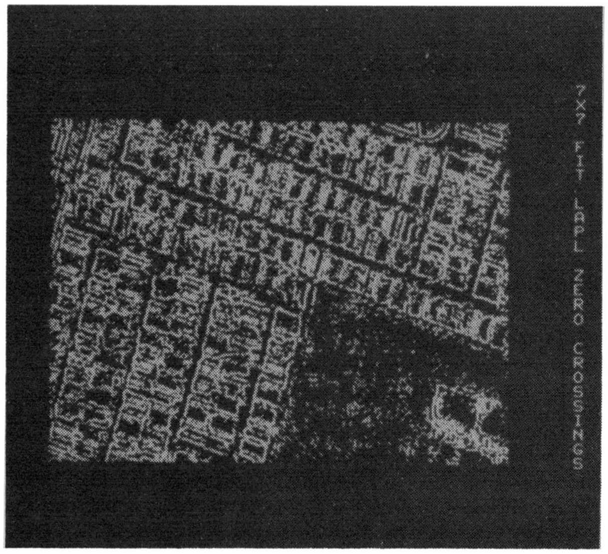

Fig. 5. The zero-crossings of a Laplacian edge detector using an equally weighted cubic fit in a 7 X 7 window. 25 percent of the pixels are assigned as edges.

is declared as an edge pixel if in the gradient direction a negatively sloped zero-crossing of the second directional derivative occurs within a distance of 0.85 of the center of the pixel and the gradient magnitude is greater than 12. The resulting image has 25 percent of the pixels assigned as edges and is shown in Fig. 4. The results are not as noisy as the Laplacian of Fig. 3. The edges are placed accurately, and they tend to be connected.

We tried an interesting variation in which we used the fitting coefficients from the bivariate cubic fit to estimate the Laplacian. The resulting zero-crossings are shown in Fig. 5, in which the zero-crossing threshold is chosen so that 25 percent of the pixels are assigned as edges. They appear more connected than the zero-crossings of the Laplacian of a Gaussian operator.

III. Discussion

There are some interesting issues which have not yet been fully discussed or understood. Whether the edge operator is a Laplacian zero-crossing one or a second directional derivative zero-crossing one, the operator must estimate partial derivatives

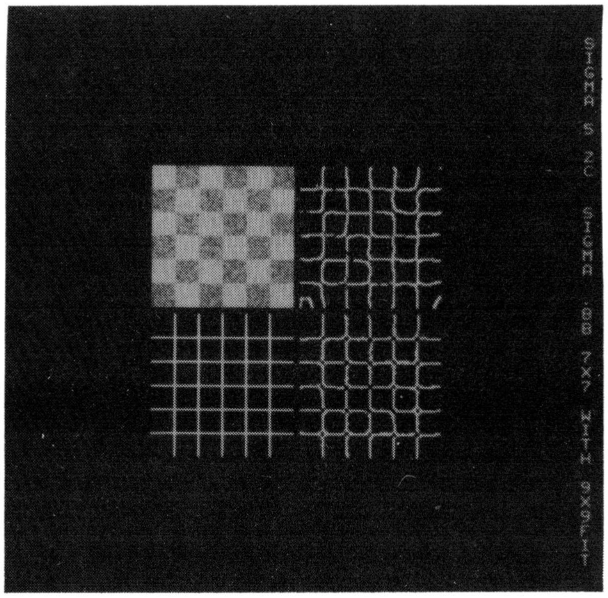

Fig. 6. The checkerboard test image (upper left-hand side), the true edge image (lower left-hand side), the zero-crossing of the Laplacian image using a Gaussian standard deviation of 5.0 (upper right-hand side), and the second directional derivative edge operator with a Gaussian presmoother having standard deviation 0.88, followed by an equally weighted cubic fit in a 9 × 9 window (lower right-hand side).

up through the third order if a zero-crossing slope is used. For a fixed neighborhood size, what is the most effective way to estimate these partial derivatives? The Marr and Hildreth scheme is equivalent to averaging and then taking finite differences to compute the partial derivatives. The Haralick scheme performs a least squares estimate assuming a local cubic polynomial model. Finite differences and least squares yield the same result only when the polynomial model has as many parameters as pixels in the neighborhood. The least squares estimate can be generalized to a weighted least square (Hashimoto and Sklansky [4] have already suggested a binomial weighted least square), and it is possible to presmooth followed by a least squares estimate. It is also possible to pose the estimation problem as a robust estimation problem, which in effect makes the weights used in the least squares fit adaptive.

We tried an example of presmoothing with a Gaussian filter having a standard deviation of 0.88 followed by a 9 × 9 equally weighted fit. Fig. 6 shows the checkerboard test image; the perfect edge image; the zero-crossings of the Laplacian of a Gaussian with a 5.0 standard deviation (upper right-hand side) result, which does not replicate the stated accuracy of the Gaussian–Hildreth experiment; and the zero-crossings of the second directional derivative edge detector (lower right-hand side). For the directional derivative edge operator, 0.8391 is the probability of a pixel being a true edge pixel given that it is assigned an edge pixel. The probability of a pixel being assigned an edge, given that it is a true edge, is also equal to 0.8391.

The Marr–Hildreth scheme chooses a direction which maximizes the zero-crossing slope of the Laplacian. The Haralick[1] and Canny [2] schemes choose the gradient direction, although they compute it in a different way. Are there other reasonable directional choices or computational techniques? What kind of experiment could be done to evaluate which is the better choice? What kind of analysis could be done to evaluate the choices in a theoretical way?

Both techniques cause edges to be displaced under certain conditions. In regions of nonlinear gray tone intensity surface, the Laplacian technique can spatially displace edges by as much as the standard deviation of the Gaussian smoother; it can even miss edges also (Berzins [1], Leclerc and Zucker [5]). Edges which curve rapidly around corners can be displaced by both techniques. There are difficulties around saddle points, especially in the second directional derivative technique which requires a nonzero gradient.

These sorts of issues and problems need to be addressed. Perhaps there could be a reader's forum on this to help us all understand the most effective way to think about the problem. Write up your idea and submit it as a note or reply to this correspondence.

REFERENCES

[1] V. Berzins, "Accuracy of Laplacian edge detectors," *Comput. Vision, Graphics, Image Processing*, vol. 27, pp. 195–210, Aug. 1984.

[2] J. F. Canny, "Finding edges and lines in images," Artif. Intell. Lab., Mass. Inst. Technol., Cambridge, MA, Tech. Rep. AI-TR-720, June 1983.

[3] W. E. L. Grimson and E. C. Hildreth, "A note on Haralick's evaluation of the Marr–Hildreth operator," *IEEE Trans. Pattern Anal. Mach. Intell.*, to be published.

[4] M. Hashimoto and J. Sklansky, "Edge detection by estimation of multiple order derivatives," in *Proc. Comput. Vision Pattern Recog. Conf.*, Washington, DC, June 19–23, 1983, pp. 318–325.

[5] Y. Leclerc and S. W. Zucker, "The local structure of image discontinuities in one dimension," in *Proc. 7th Int. Conf. Pattern Recog.*, Montreal, P.Q., Canada, July 30–Aug. 2, 1984, pp. 46–48.