DETECTION OF EDGES USING RANGE INFORMATION

Amar Mitiche[2]
J. K. Aggarwal[1,2]

TR-82-4 October 1982

[1]Department of Electrical Engineering
[2]Laboratory for Image and Signal Analysis
The University of Texas at Austin
Austin, Texas 78712

This research was supported by the Air Force Office of Scientific Research under grant AFOSR 82-0064.

ABSTRACT

Range data provide an important source of 3-D shape information. This information can be used to extract jump boundaries which correspond to occluding boundaries of objects in a scene and "edges" which correspond to points lying between significantly different regions on the surface of objects. We are mainly interested in range data obtained from sensors such as lasers. The main problem with this type of range finders is the fact that the accuracy of the measurements depends on the power of the signal that reaches the receiver. This study describes how a range edge detection procedure can be designed that has low sensitivity to noise and imbeds all the knowledge available on the range measurement accuracy.

1. Introduction

Range data provide an important source of 3-D shape information. Range data implicitly contain information about the shape of the surface of objects because the coordinates of points on the surface of these objects can be easily recovered from them. Moreover, they can be used to extract jump boundaries which correspond to occluding boundaries of objects in a scene, and "edges" which are points that lie on the intersection of two regions on the surface with significantly different parameters (e.g. the edge between two visible faces of a cube). Jump boundaries and edges are important cues in the segmentation process because they delineate the extent of surfaces. These boundaries and edges are intrinsic properties of the surface of objects unlike edges in "intensity" images derived from range data. Lynch ([1]) for example creates an image by using a range coding convention in which dark = near and light = far. The image is then enhanced to compensate for low frequency trends present in systems which have shallow line of sight. Finally, edges are computed using the Sobel operator. In this study we are concerned with edges that actually occur on the surface of objects in the scene. Although simple methods have been devised [2] that are successful at computing jump boundaries, the problem of finding edges is a much more delicate problem which has not received enough attention [3,4,5]. Here, we are mainly interested in range data obtained from sensors such as the one described in [6]. The major problem with this type of range finder is that the accuracy of the measurements depends on the power of the signal that reaches the receiver. The accuracy of the range data is therefore dependent on the parameters of the system (transmitted beam power and receiver variables) and on the characteristics of the target (orientation and reflectance of its surface and actual distance from the sensor). System parameters and target characteristics affect accuracy because they affect the strength of the signal that returns to the sensor. It should be pointed out that we do not restrict ourselves to the SRI laser sensor [6]. But, the well documented research described in [2] and [6] provides a good example of the current activities in the field.

Our goal in this study is to design and analyze a procedure for detecting edges using range information that has low sensitivity to noise. We also want to relate the range measurement accuracy to the problem of detecting edges. The input to this procedure includes range data and a model for range measurement error. Basically, this procedure determines the best partition of a neighborhood of each point in the scene into two contiguous regions. Planes are fitted to these regions and a measure of the goodness of fit is calculated. The measure should imbed all the knowledge that one has about range measurement accuracy, so that the resulting value can be used to select the "best" partition. Subsequent analysis can then extract significant edges from the scene. Although we are fitting planes to small patches on the surface of objects, we are not looking for planar surfaces as in [3]. We are interested in determining the presence or absence of an edge at points on the surface of objects.

The remainder of this paper is organized as follows: Section 2 is a formal definition of the problem, Section 3 is a description of the criterion for selecting the best partition and Section 4 contains an overall description of the computational procedure for edge detection and presents some experimental results. Finally, Section 5 contains a summary.

2. Problem Definition

Let X be a point in the range image. Here range image refers to the set of 3-D points obtained from the range finder. The problem is to determine whether an edge is present at X. Let N be an appropriately sized neighborhood of X and let N be divided into two contiguous regions C_1 and C_2. Here, two range points are neighbors if they project onto adjacent points in the focal plane. The line connecting a point X in the range image to the view point (sensor) is the projecting ray and intersects the focal plane at X' as in Figure 1. Thus a neighborhood of X' defines a neighborhood of X. Let the neighborhoods in the focal plane corresponding to regions C_1 and C_2 be called C'_1 and C'_2. With the assumption that C_1 and C_2 are planar surfaces, let n_1 and n_2 be the surface normals of the planes best fit to the points in C_1 and C_2 respectively. The planes are best fitted to the regions according to some selected criterion such as the least squared error criterion or some other functional optimization procedure. The collection of points in N together with these best fit planes will be called a partition of N and denoted

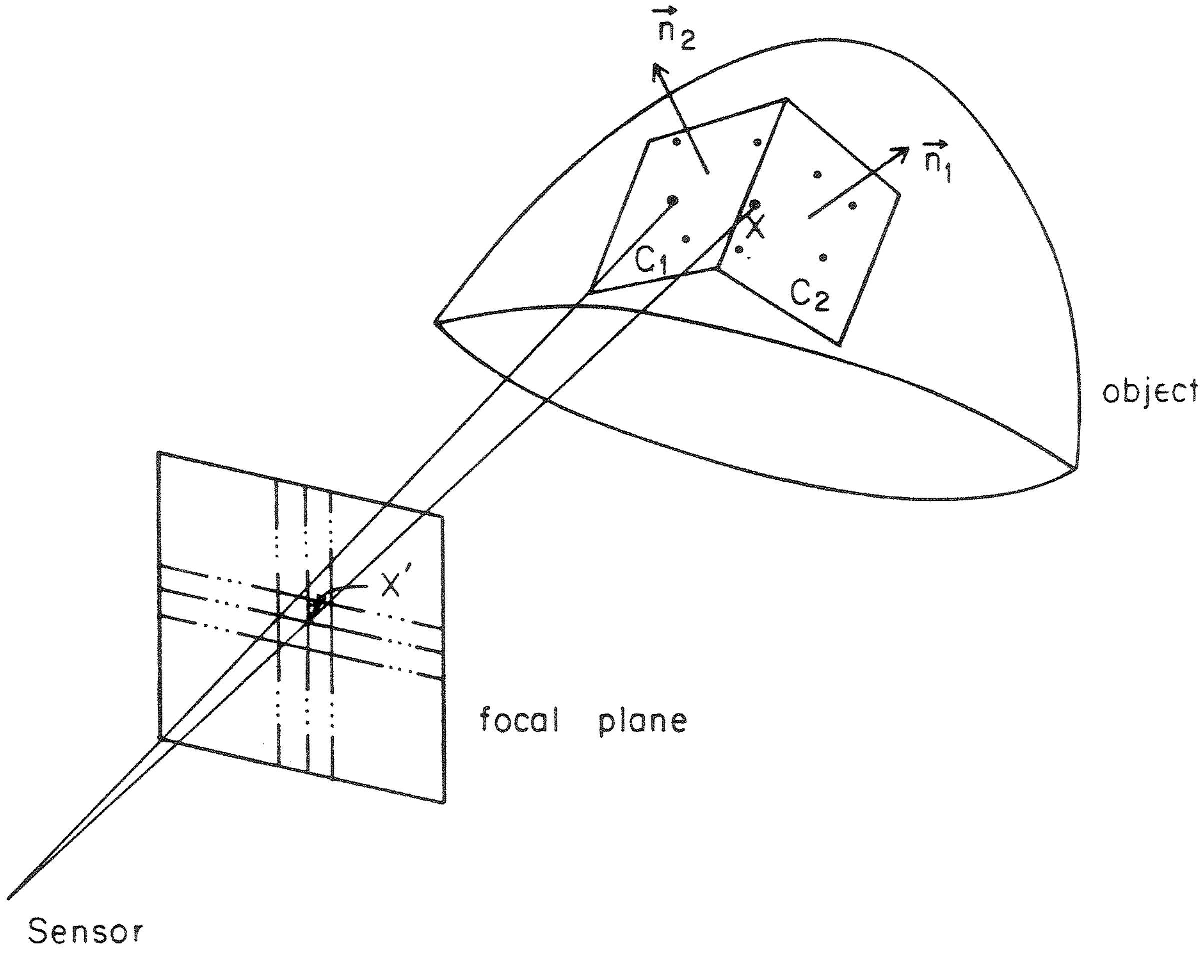

Figure 1. A partition at point X.

P. An example of this situation is shown in Figure 1. Finally let ϵ be a measure of the quality of fit of these planes. For example, this measure could be the minimum error associated with fitting the planes to the observed points in N, if these are fitted according to the least squared error criterion. But, for sensors such as lasers, ϵ should relate to the surface normals, n_1 and n_2, the surface reflectance of C_1 and C_2 (we assume that the reflectance is essentially constant over each of these regions) and the actual distance of the points in N to the sensor, since these parameters influence the accuracy of range measurements. In general, ϵ is a measure computed according to some model for range measurement errors and should reflect how likely the points in regions C_1 and C_2 lie on planar surfaces. An example illustrating this idea is reported in Section 4.

The problem of finding edges in the range image can now be formulated in its generality as follows:

"Find the best partition $P(n_1, n_2, \epsilon, f)$ from a set of partitions of N and test the significance of the difference between n_1 and n_2 to declare the presence or absence of an edge at X."

Here f is the intensity function. The intensity function can be used in conjunction with range data to enhance the edge detection performance [5]. This aspect of the problem will, nevertheless, be viewed as a separate process that will not be the subject of the present analysis.

3. Range Edges

The problem as defined above may suggest an analogue to the sloped facet model of Haralick [7] for edge and region analysis in intensity pictures. Using this model, an edge at point X is computed by using an F-statistic to test the significance of the difference between the slopes of the planes best-fitted to appropriate neighborhoods of X. The problem considered here is much different in the sense that we are not dealing with simple function graphs but with actual object surfaces. In the sloped facet model, an ideal region is a sloped plane and the observed region is obtained from the ideal region by adding random noise to the intensities (z-values) in that region. Thus, in terms of the (x,y,z)-coordinates of the resulting points, only the z-coordinate is a random variable. In the problem considered here, the x-,y- and z-coordinate of points on the surface of objects should all be considered random (in the sense that random noise is present). A straightforward statistical analysis as in [7] is not possible.

Instead, a Baysian approach is taken here [8]. Given range measurement as evidence, we want to know how likely a given partition is. In other words, we want to know how good a model the planes in a given partition are for the observed points of this partition. Consider a set of partitions $\{P_i\} i=1,...n$ of an appropriately sized neighborhood of point X as before and let S be the sample of points in this neighborhood. Then the conditional probability of partition P_i given the sample S can be written as:

$$\mathrm{Prob}(\Gamma_i \mid S) = \frac{p(S \mid \Gamma_i) \cdot \mathrm{Prob}(\Gamma_i)}{p(S)} \qquad i=1,\ldots n \qquad (1)$$

where $p(S) = \sum_{j=1}^{n} p(S \mid \Gamma_j) \cdot \mathrm{Prob}(\Gamma_j)$

Here $p(S \mid \Gamma_j)$ denotes the partition- conditional probability density function of the sample S and $\mathrm{Prob}(\Gamma_j)$ is the a-priori probability of partition Γ_j. If the edges in the scene under consideration are known to have high directionality then the a-priori probabilities could be biased and a high value assigned to a subset of them. The denominator in the right hand side of (1) is the same for all i and if the prior probabilities $\mathrm{Prob}(\Gamma_i)$ are all equal, the problem of selecting the best partition reduces to computing $p(S \mid \Gamma_i)$ for each partition Γ_i and selecting the one that yields the highest value. In this study we will consider all partitions equally likely.

Since these partition-conditional probabilities are the basis for comparing partitions, we want to imbed in them any knowledge we have about the accuracy of range measurements. In general, given the variables of the sensing instrument and the properties of the target under some basic assumptions, some knowledge can be derived about the distribution of the range measurements. For example, analytic expressions are derived in [5] that relate the mean and variance of range data to the physical properties of the target and the parameters of the sensor.

Let X be a point on the target surface and let r_X be the random variable that describes the range measurement at X, i.e., the measured distance from the sensor location to the point X. If we let $p(r_X|\Gamma_i)$ be the conditional probability density function of the range measurements given partition Γ_i (i.e. given that the point under consideration actually lies on one of the planes of the partition Γ_i) and assuming independence of these measurements from point to point we can write:

$$p(S|\Gamma_j) = \prod_{X \in S} p(r_X|\Gamma_j) \quad j=1,n$$

As mentioned earlier r_X is a function of the physical properties of the target (surface orientation with respect to the signal beam direction, reflectance and actual distance from the sensor) and sensor parameters (fixed for a given scene). For example, given a partition Γ_j, r_X can be assumed to be normally distributed with mean $\mu_X = R_X$ and standard deviation σ_X, where R_X is the actual distance to point X, i.e., the distance from the sensor locus to the projection of X onto the appropriate plane of the partition Γ_j along the direction of the signal beam. Also $\sigma_X = \sigma_X(R_X, \rho, \Phi)$ where R_X is defined as before, ρ is the surface reflectance (considered constant over the neighborhood in question) and Φ is the orientation of the appropriate plane in Γ_j with respect to the signal beam direction (See Figure 2). In the example reported in this study, r_x is normally distributed with mean R_x and constant standard deviation.

Thus we have a method of selecting the best partition of a neighborhood of a point that will imbed all the knowledge we may

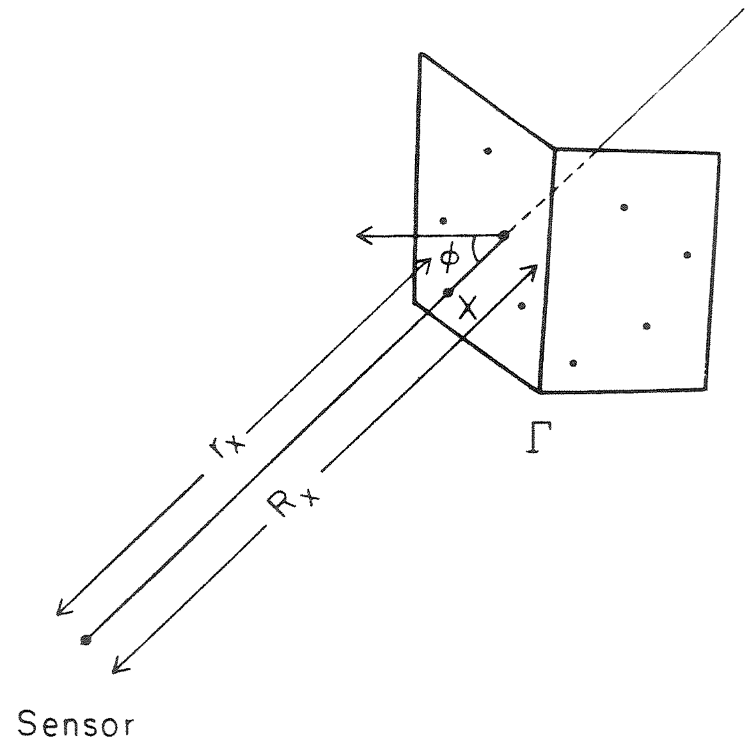

Figure 2. A model for range error.

have about the range measurement accuracy. What we end up with is a map of angles. Each angle is the angle between the normals of the best fit planes of the most likely partition of an appropriately sized neighborhood of the point at which it has been computed. Moreover, a probability is assigned to each of the angles and if an edge is present at a point, an orientation can be assigned to it which is the orientation of the intersection of the two planes of the best partition at this point. Although the normals to the best fit planes contain more information than just the value of the angle between them, only this value is of interest to us in this study. It is used, along with other evidence, to determine the likelihood of edge presence. For display purposes, we will be concerned only with the orientation in the focal plane of the line between the neighborhoods C'_1 and C'_2 (in the focal plane) corresponding to the best partition at this point. This orientation will be referred to as the orientation of the edge. In the above analysis the intensity picture has not been used in the selection of the best partition. The cooperation between, or combination of, intensity and range data can be exploited as an enhancement process that would hopefully yield a description of the target scene better than the description which could be obtained by either intensity or range alone. This aspect of the problem is discussed in [4] and will not be considered here.

4. Computational Procedure

The computational procedure contains basically four steps:

extraction of jump boundaries, computation of partitions, dismissal of flat surfaces and a non-maxima suppression step to select the significant edges. These steps are described below in more detail.

(a) <u>Extraction of Jump Boundaries</u>

Jump boundaries occur at points where there is a significant change (jump) in the range values. These boundaries can then be detected by comparing range values at neighboring points. This procedure may, however, falsely detect jump boundary locations at points on highly oblique surfaces such as a plane the orientation of which is nearly orthogonal to the line of sight. A better method would look for significant jumps in the first order differences of range values at neighboring points. Non-maxima suppression based on these differences can be used to thin out these boundaries.

The subsequent steps of the procedure are aimed at detecting edges at points other than jump boundary locations.

(b) <u>Computation of Partitions</u>

Let N designate the appropriate neighborhood of a point in the range image and let n be the number of directions in which edges are to be computed. For example, n could be 4 and the directions are horizontal, vertical, left diagonal and right diagonal. For each direction planes are fit according to some selected criterion to the corresponding two regions in the neighborhood N, thus determining the partitions $\Gamma_i, i=1...,n$. The angles associated with these partitions are denoted by $\theta_i, i=1,...,n$.

(c) Dismissal of Flat Surfaces

The elimination of points on flat regions can be accomplished in two steps.

1) Discard all points for which $\theta_i < t$ $i=1,\ldots,n$. This thresholding step is intended to eliminate deep "interior" points before computing the likelihood of the partition at these points. Deep interior points are surface points far from edges or jump boundaries of surfaces. This step saves a significant amount of computation time since calculating the likelihood of a partition might be a costly operation.

2) Points that survive the preceding thresholding operation are considered next. The likelihood of each of the partitions at these points is computed and the most probable partition, $\hat{\Gamma}$, is selected. This likelihood is determined as described in Section 3 by computing the distances from the points contained in the partition to the appropriate plane and using the selected error or range accuracy model. Now discard all points for which $\theta < t$ where θ is the angle associated with the best partition $\hat{\Gamma}$.

(d) Non-Maxima Suppression

At this stage, the remaining points are potential edge points. These points are generally clustered around true edge position because of the overlap of neighborhoods at adjacent or nearby points. If Prob $[\hat{\Gamma}]$ denotes the probability of the most probable partition $\hat{\Gamma}$, then this step will discard all points for which Prob $[\hat{\Gamma}]$ is not a local maximum in an interval orthogonal to the edge. Non-maxima suppression is intended to discriminate

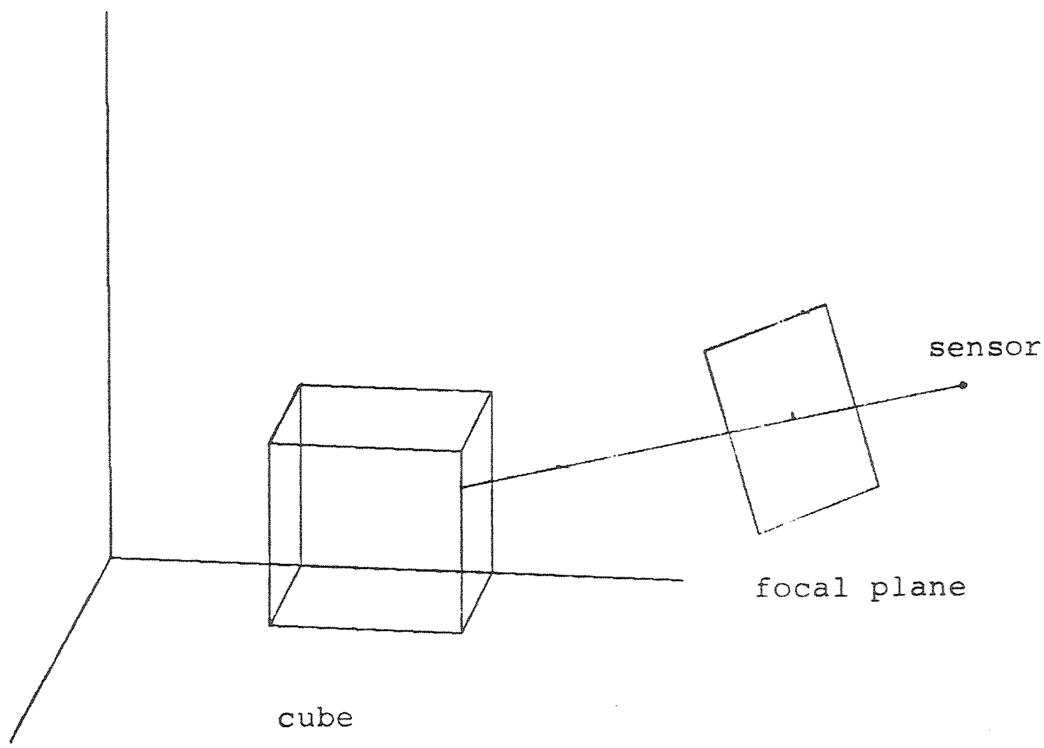

Figure 3. The ideal cube scene, focal plane and sensor.

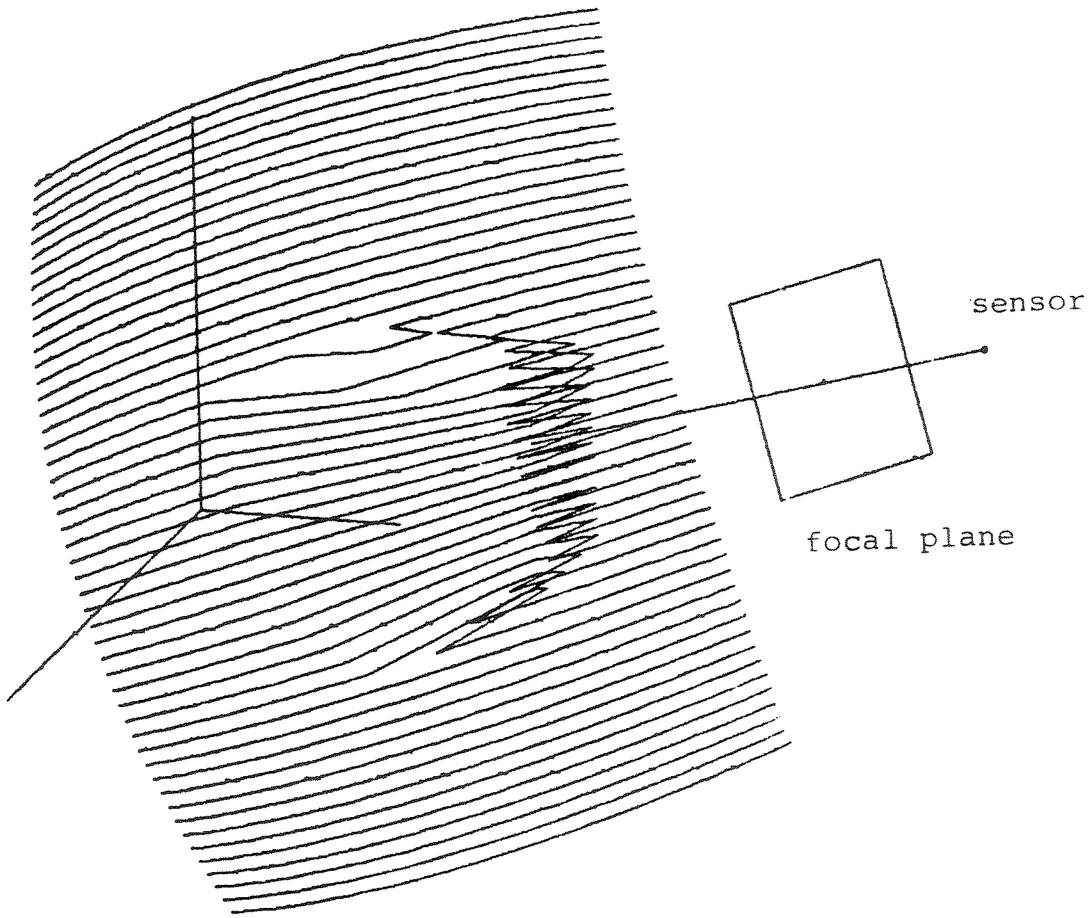

Figure 4. A line drawing perspective view of the ideal cube scene of Figure 3 obtained from range data by joining points when the scene is scanned horizontally in raster fashion.

5. It does not depend on a particular sensor.

In our experiments, edges are detected in the horizontal and vertical directions (n=2) using 5x8 and 8x5 neighborhoods respectively, which seem adequate for our example. Determining an optimal size neighborhood for a given class of object models is a hard and interesting problem that we have not considered here. Some effort has been devoted to this problem by Davis and Mitiche [9] in the context of intensity images. Shaping neighborhoods differently to look for edges in different directions seems to give better results than keeping the same neighborhood shape for all edge directions. The best fit criterion is the least squared error criterion. This implicitly assumes that the model for measurement errors does favor small errors in the measurement of range. For a model that does not favor small errors, the probability density function that describes it should be normalized to have zero mean and the range data should be modified accordingly. The angle threshold, t, is 30^o and the interval in which non-maxima suppression is performed is of width 5.

Noise was added to the range of the ideal cube scene. This noise is normally distributed with zero mean and standard deviation equal to 0.1. This represents an appreciable amount of noise since approximately 5% of the errors are greater than one-fifth the side length of the cube and 32% of them are greater than one tenth of it. This noise is considerably higher than the noise for the office scene reported in [5] where distances were sampled enough times to bring the accuracy within a couple of centimeters. For relatively noisier scenes, larger neighborhoods are necessary. The computed edges before non-maxima suppression of step (C) are shown in Figure 5. Note the clusters of responses. The final edge map for this noisy cube scene is shown in Figure 6. Due to the fact that we are looking for edges in the horizontal and vertical directions, the detection of edges which actually lie in other directions is not accurate. But basically all the edges in the cube have been extracted although some noise edges are present. The edges on the left of the cube in Figure 6 are noise created edges that could not be dismissed by the computational procedure described in the paper. A possible remedy to eliminate them would be to consider the probability of the edge presence assertion at these points. But we chose not to do this to provide a fair comparison with the results depicted in Figure 7 and Figure 8. For comparison, the edge detection procedure was used on the same noisy cube scene but with the least mean squared error of fit as the criterion for selecting the best partition. Results are shown in Figure 7. Note that several edges have been falsely dismissed. Minor isolated noise was removed before edge maps are displayed in Figures 6 and 7.

Finally, and again for comparison, a simple edge detection procedure based on pointwise computation of angles, thresholding and non-maxima suppression as described in [4] yields the edge map of Figure 8. Jump boundaries were extracted by looking at the difference in range at neighboring points. This method is obviously very sensitive to noise and indicates the need for larger neighborhoods to reduce the effect of noise. Figure 9 shows that only marginal improvement is obtained when the range

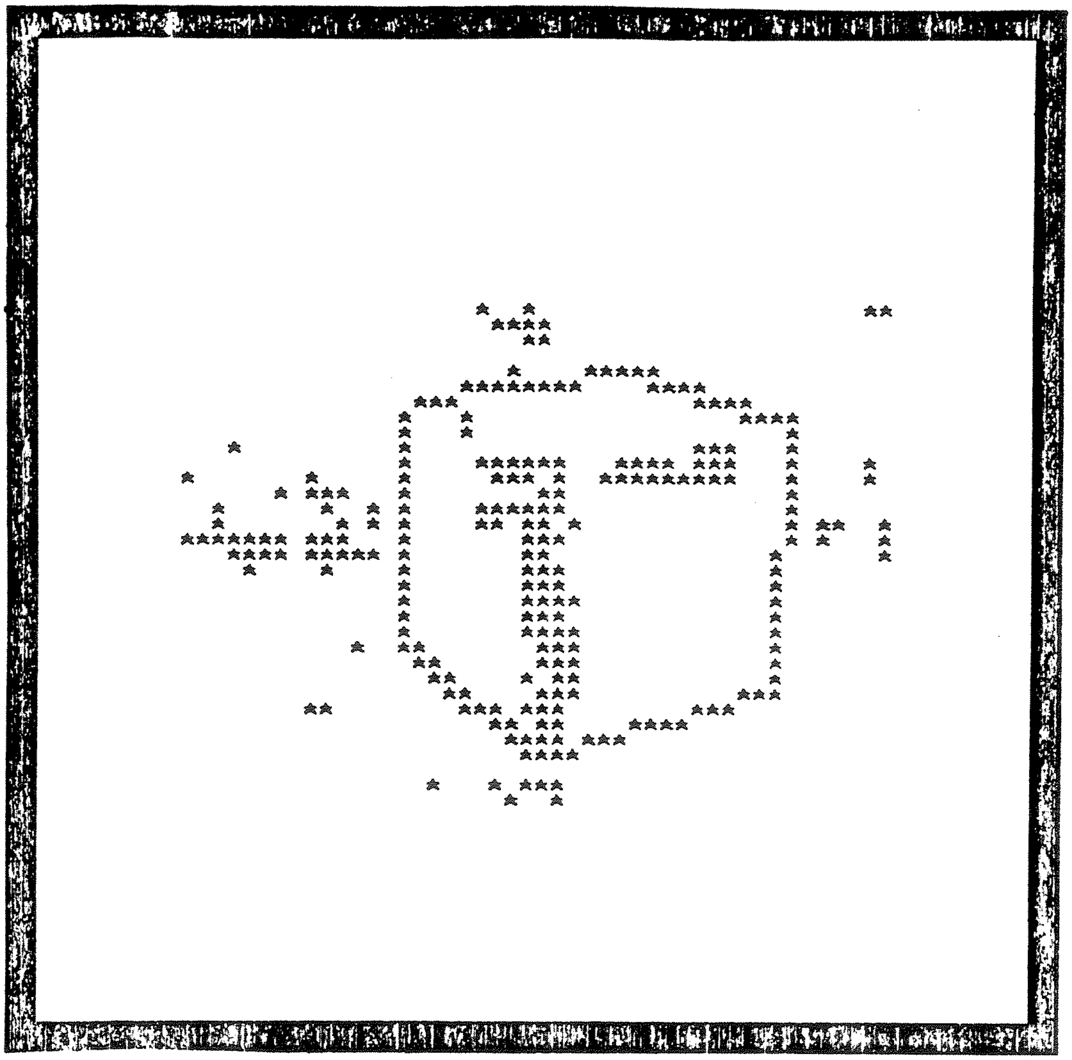

Figure 5. Probabilistic model - Edge map before non-maxima suppression.

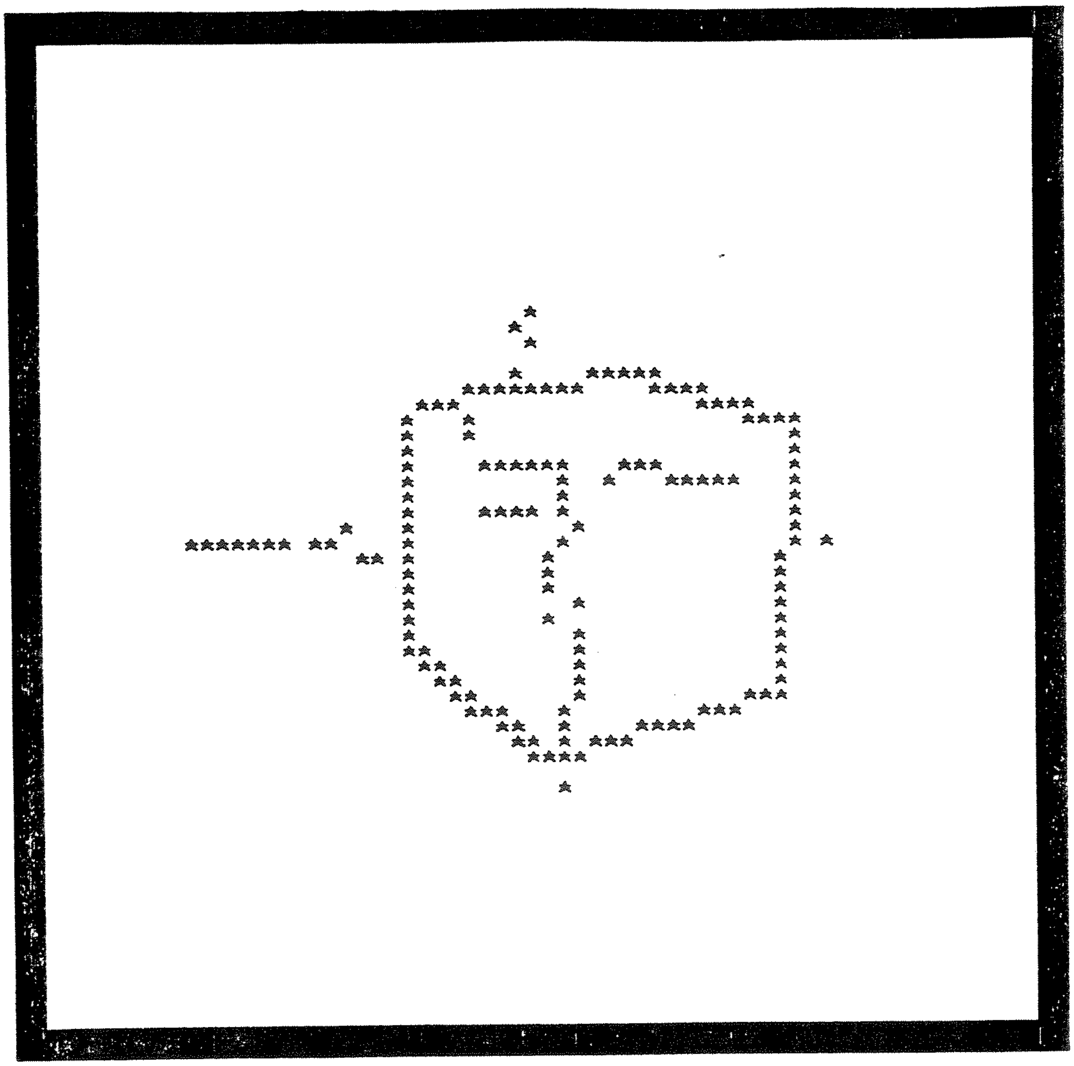

Figure 6. Probabilistic model - Final edge map.

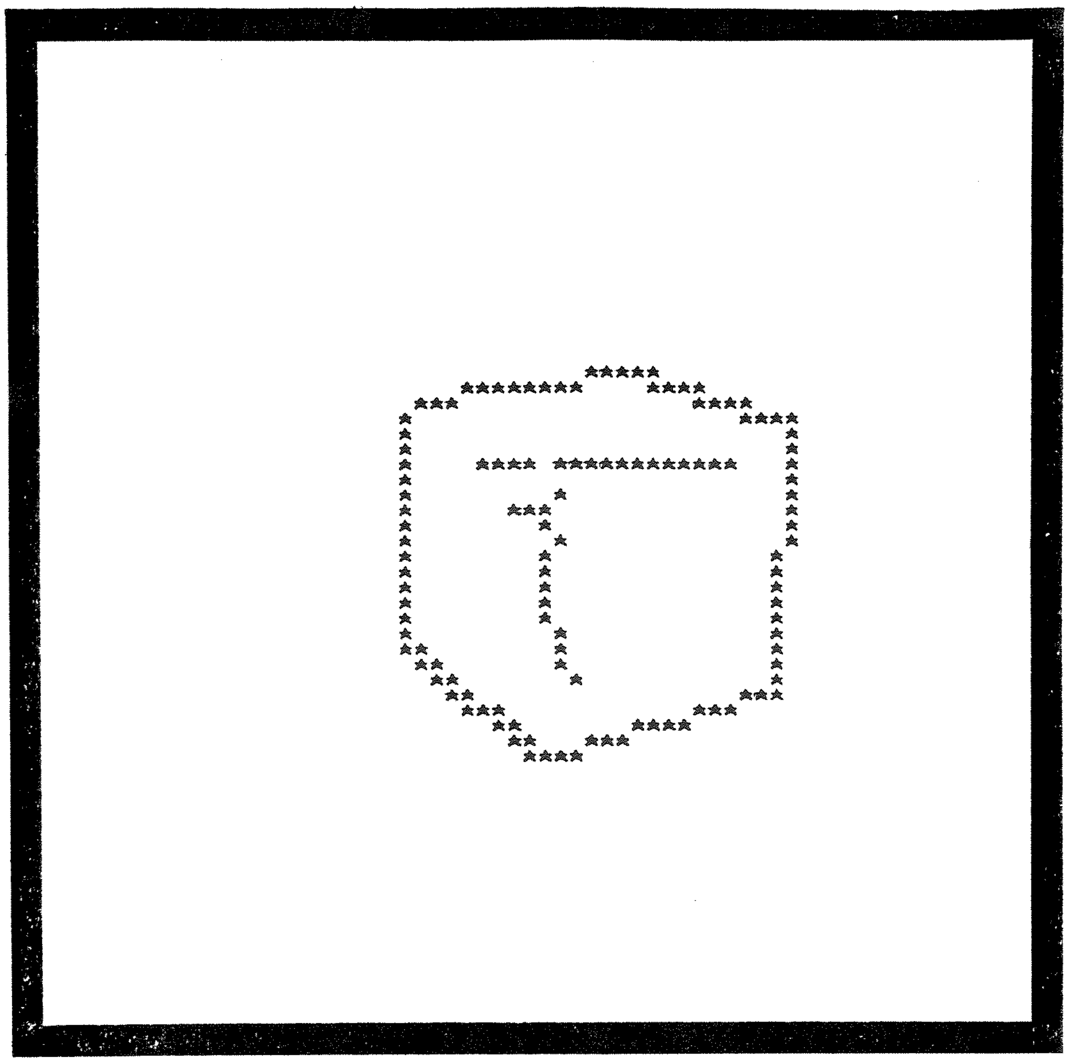

Figure 7. Least squared error model - Final edge map.

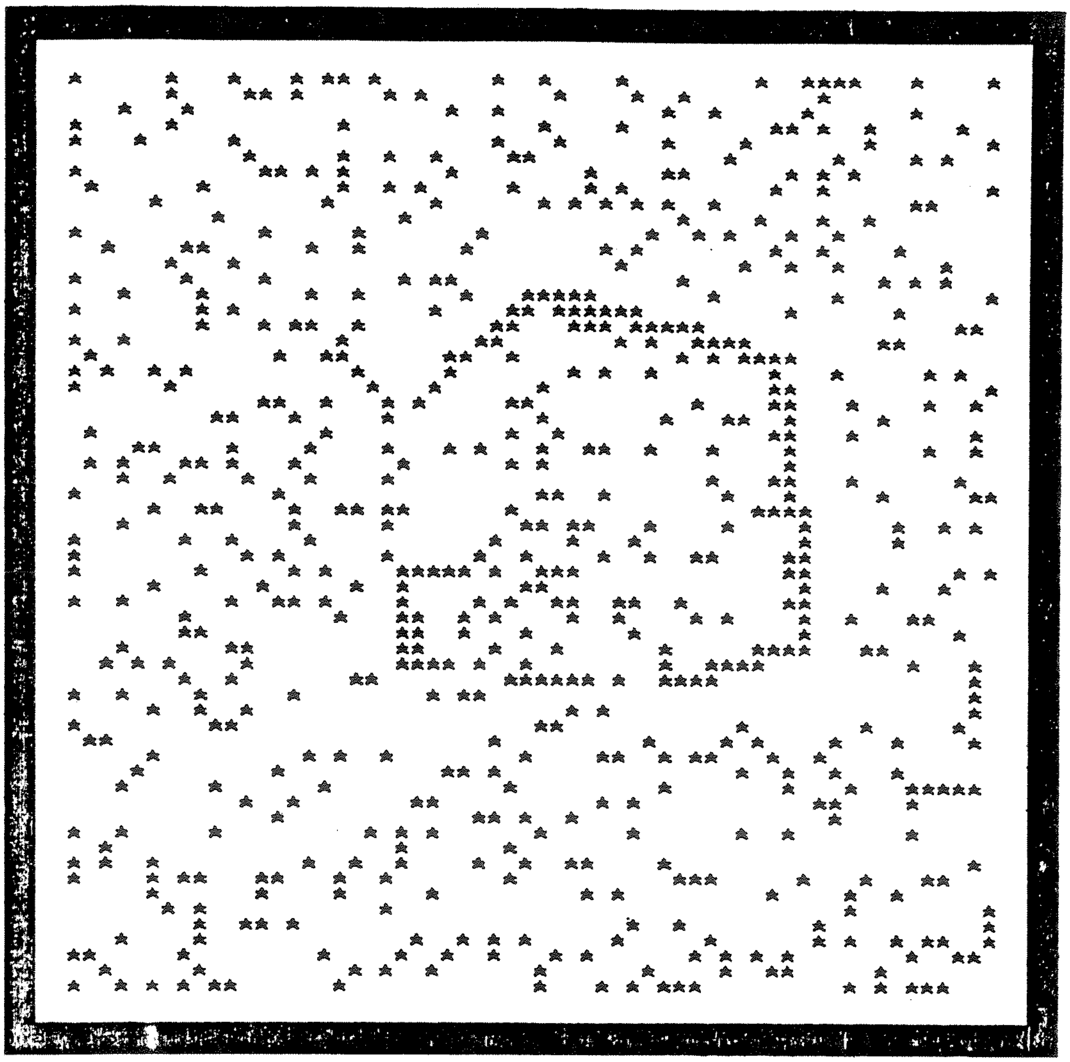

Figure 8. Pointwise computation - Final edge map.

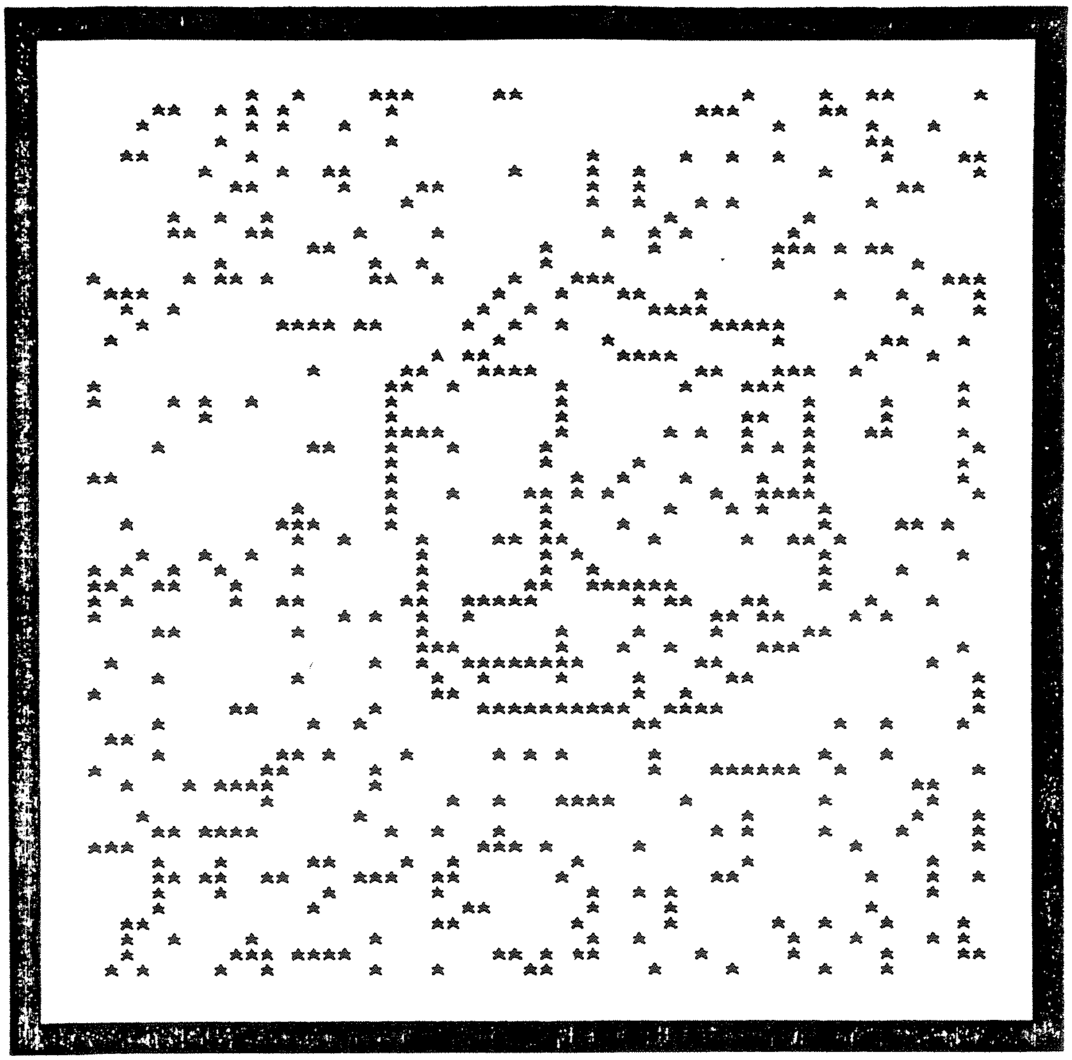

Figure 9. Pointwise computation - Final edge map using smoothed range.

data is smoothed. Performance with this method is, nevertheless, good on almost perfect data.

5. Conclusion

This paper has discussed the problem of designing a method for detecting edges in a scene using range information. Reliability was the guiding requirement in the design of this procedure. The method proposed here is based on determining the best partition of an appropriate neighborhood of a point into two contiguous regions assumed to be planar surfaces.

It was shown that local computation of edges using larger neighborhoods performed sensibly better than pointwise computation of edges which is very sensitive to noise. Moreover, it was shown that it is possible to design a procedure that incorporates all the available knowledge about range measurement accuracy, resulting in a better edge detection performance.

Most of the computational time (our example runs in about 2 minutes on a PDP-10) is spent on computing the likelihood of partitions for potential edge points. The problem is then basically linear in the number of partition edge directions one is willing to consider. When the computation of likelihoods contain time consuming function calls, efficiency can be improved using a table to organize the function values and table lookup to replace function calls. In this paper, range edges were detected in the horizontal and vertical directions. The partition selection process to best represent a neighborhood in the range image can benefit from the use of more partition edge directions. These can be straight line directions or curved partition boundaries. But in this case, it is not quite clear how non-maxima suppression should be performed.

Results presented herein indicate that it is possible to detect edges reliably enough in the presence of an important amount of noise. Thus, repeated and time consuming sampling of range to bring accuracy of measurements within a given tolerance as reported in [6] may not be necessary.

References

1. D. K. Lynch, "Range Image Enhancement via One-dimensional Spatial Filtering," Computer Graphics and Image Processing 15, pp. 194-200, 1981.

2. R. O. Duda and D. Nitzan, "Low-Level Processing of Registered Intensity and Range Data" in Proc. 3rd International Joint Conference on Pattern Recognition, Coronado, CA, pp. 598-601, Nov. 1976.

3. S. Inokuchi and R. Nevatia, "Boundary Detection in Range Pictures," Proc. 5th Int. Conf. on Pattern Recognition, Miami Beach, FL, pp.1301-1303, Dec. 1980.

4. D. L. Milgram and C. M. Bjorklund, "Range Image Processing: Planar Surface Extraction," Proc. 5th Int. Conf. on Pattern Recognition, Miami Beach, FL, pp. 912-919, Dec. 1980.

5. B. Gil, A. Mitiche and J. K. Aggarwal, "Experiments in Combining Intensity and Range Edge Maps," Computer Graphics and Image Processing, (to appear).

6. D. Nitzan, A. E. Brain, R. O. Duda, "The Measurement and Use of Registered Reflectance and Range Data in Scene Analysis," Proc. IEEE 65, pp. 206-220, Feb. 1977.

7. R. M. Haralick, "Edge and Region Analysis for Digital Image Data," Computer Graphics and Image Processing 12, pp. 60-73, 1980.

8. R. O. Duda and P. E. Hart, Pattern Classification and Scene Analysis, John Wiley & Sons, Inc., New York, 1975.

9. L. S. Davis and A. Mitiche, "Edge Detection in Textures," Computer Graphics and Image Processing 12, pp. 25-39, 1980.