

OBJECT DETECTION AND MEASUREMENT USING STEREO VISION*

Donald B. Gennery
Artificial Intelligence Laboratory
Computer Science Department
Stanford University
Stanford, California 94305

A method of detecting and measuring objects for the purpose of representing three-dimensional outdoor scenes, using data such as that obtained from stereo vision or from a scanning laser rangefinder, is described. Objects are approximated by ellipsoids. Segmentation of the objects from the background and from each other is done by finding the ground surface, forming a preliminary segmentation by clustering the points above the ground by more than a threshold, fitting ellipsoids to match these clusters and to avoid obscuring the other points, and adjusting the clusters according to the fits. The method is designed to produce results useful for obstacle avoidance and navigation in an exploring vehicle, such as a Mars rover. An example is given showing results obtained from a stereo pair of pictures of Mars from the Viking Lander.

1. Introduction

In a previous paper [1], a stereo vision system and its possible use in an exploring vehicle was described. This paper describes further work towards a complete system for an exploring vehicle, specifically, the use of three-dimensional data for the detection of objects and the measurement of their position, size, and approximate shape. Although the three-dimensional data could be obtained from a laser rangefinder, the object detector is designed to be tolerant of errors in this data, such as mistakes produced by incorrect matches in stereo vision data and poor accuracy of distances from stereo.

Many approaches are possible in describing the shapes of objects. For example, generalized cones [2] can be used to describe complicated objects. However, in some cases the resolution of the data is insufficient to produce detailed information about the shape. In other cases, the objects are so irregular as to make such detailed descriptions very difficult. However, in such cases information about the position and size and some crude information about the shape may still be quite useful. For example, for obstacle avoidance in a roving vehicle, this sort of information is adequate. Furthermore, this sort of information for each object in a large scene containing many objects amounts to quite detailed information concerning the whole scene, and thus it would be useful for navigation.

For these reasons, and because of its convenient mathematical properties, here each object will be approximated by an ellipsoid. By "object" we do not necessarily mean an actual physical object, but merely a portion of the scene that can be reasonably approximated by an ellipsoid. Thus, if we use as an example a vehicle exploring Mars, an object may be a single rock on the Martian surface, two or more adjacent rocks, or merely a bump in the ground. Also, an L-shaped physical object might be represented as two objects.

This ellipsoidal representation should be quite appropriate for representing rocks on Mars, because rocks probably tend to resemble more nearly ellipsoids than any other simple shape. However, it could also be used to represent cars in a parking lot or trees in a field, for example, especially in aerial photographs where the resolution may be poor compared to the size of the objects, and in other cases where precise object description or recognition is not necessary but rather an overall description of the scene is desired.

The stereo vision processing or laser rangefinder results in data representing the three-dimensional position of a large number of points distributed over the scene. The first step in the processing of this three-dimensional data is to find the ground surface. A method of doing this was previously described [1]. In general, an entire scene would be partitioned into small areas, in each of which the ground would be approximated by a plane or paraboloid. Then points which are above the ground by a sufficient amount (depending on the computed accuracy of the points, the roughness of the ground, and the minimum size of object that is of interest) are candidates for points on objects.

* This work was performed under NASA contract NASW-2916 and the Defense Advanced Research Projects Agency contract MDA903-76-C-0206.

These above-ground points are clustered to produce preliminary groupings of points which correspond roughly to objects. An ellipsoid is fit to each cluster by first computing an initial approximation based upon the moments of the points in the cluster and then iterating a weighted nonlinear least-squares adjustment to fit the ellipsoids to these points and to avoid obscuring other points. Then, according to the relative positions of the ellipsoids and points, clusters can be broken or merged, and the process repeats until the apparently best segmentation is found. Each of these steps will be described in the following sections.

The object detection and measurement process as described here uses only the three-dimensional position information. The brightness information is discarded after the stereo processing. However, a more complete system would use both types of information. Perhaps an edge detector could be applied to the brightness data in the regions near the outlines of the ellipsoids in order to refine the boundaries of the objects, for example.

Matrix notation will be used throughout this paper. Matrices will be denoted by capital letters. The transpose of a matrix A will be denoted by A^T , and the inverse of A will be denoted by A^{-1} . The trace of a square matrix A (sum of the diagonal elements, which is equal to the sum of the eigenvalues) will be denoted by $\text{tr}(A)$. A vector in three-dimensional space will be represented by a 3-by-1 matrix containing the Cartesian coordinates, usually denoted by X with an appropriate subscript. (Hohn [3] provides a good text on matrix algebra.)

2. Preliminary Clustering

Once the ground surface has been determined, all points that are above this surface by more than a threshold are clustered to form an initial approximation to the segmentation of the scene into objects.

Various clustering techniques could be used here. One possibility is a relaxation method, such as Zucker's [4]. However, at present, the clustering is done by using the minimal spanning tree of the points. (The minimal spanning tree is the tree connecting all of the points such that the sum of the edges is minimum.) This is computed by using the nearest neighbor algorithm [5]. (The length of the edges of the tree is defined here as the three-dimensional Euclidean distance between the points.) Then the tree is broken at every edge whose length is greater than twice the average length of the adjacent edges [5]. However, a minimum length for an edge to be broken (related to the resolution of the data) is specified, so that the method will not be overly sensitive to local fluctuations in the data. Also, a maximum can be specified, beyond which all edges are broken.

3. Initial Approximations to Ellipsoids

Since each ellipsoid will be fit to a cluster of points by an

iterative process, an initial approximation is needed. A good approximation increases the likelihood of convergence, decreases the number of iterations required, and can be used as the result in case the iterations do not converge. This initial approximation is obtained from the three-dimensional moments, through the second order, of the points in the cluster.

An ellipsoid can be represented by the following matrix equation:

$$(X - X_c)^T W (X - X_c) = 1$$

where X is a vector of the three-dimensional coordinates of any point on the surface of the ellipsoid, X_c similarly is the position of the center of the ellipsoid, and W is a positive-definite symmetrical 3-by-3 matrix. (See, for example, Hohn [3].) Let M denote the inverse of W . (The square roots of the eigenvalues of M are the lengths of the semi-axes of the ellipsoid.) The relationship between the computed moments and the matrices X_c and M depends on the distribution of points over the ellipsoid. If the object has been viewed from a single point, we will have points distributed nonuniformly over half of the surface. (Actually slightly less than half will be seen because of perspective. Also, in stereo vision, both cameras must see each point, so that with a single pair of cameras only the common area seen from both camera positions will appear. These two effects will be neglected below, however.)

We assume here that the object is seen from a single viewpoint by a raster scanning device which produces points distributed uniformly in the image plane. Such a device might be a scanning laser rangefinder or an area-based stereo system. Actually, because of missing points, the distribution will not be uniform. It would be possible to estimate the actual distribution by computing higher-order moments, but this might be overly sensitive to randomness in the distribution or an inadequate density of points, so it is not attempted here. As an approximation, we assume an orthogonal projection instead of a central projection. Let X_g denote the vector of normalized first moments (centroid) and M_g denote the matrix of second moments about X_g obtained with this distribution, and let X_0 denote the position of the camera.

The relationship connecting X_g and M_g to X_c and M can be derived by first considering the case of a sphere of radius r . A little integration shows that in this case the eigenvalue of M_g corresponding to the eigenvector $X_0 - X_c$ is $r^2/18$, the other two eigenvalues are both $r^2/4$, and X_g is X_c plus $2r/3$ times the unit vector in the $X_0 - X_c$ direction. All three eigenvalues of M should be r^2 in this case. An ellipsoid can be considered to be a distorted sphere (using stretching and skew distortions). Thus the ellipsoid can be considered to be stretched in the various directions by the amount given by the square roots of the ratios of the above eigenvalues, but in computing the displacement of the center, instead of r the distance from X_c towards X_0 to the ellipsoid surface must be used. Thus the displacement of the center is the vector

$2(X_o - X_c)/3$ divided by the scalar $\sqrt{(X_o - X_c)^T W (X_o - X_c)}$. Since X_c , X_s , and X_o are colinear, X_c can be replaced by X_s without changing the value of this ratio. Also, $W (=M^{-1})$ can be replaced by $M_s^{-1}/18$ in this expression, because of the stretching discussed above. Thus X_c can be computed from X_s by translating by this amount. To compute M , we can take 4 times M_s to account for the factor of 4 in two dimensions, but this leaves 14 out of the factor of 18 by which we need to stretch the moments in the direction toward the camera. This extra amount can be added by adding 14 times the moment produced by a fictitious point at the intersection of the X_s -to- X_o line and the surface of the ellipsoid corresponding to M_s . In order to keep the ellipsoid to a reasonable shape when there are not enough points to determine it well, M as obtained above is averaged with a scalar matrix whose diagonal elements are s (which represents a sphere of radius $\sqrt{s}$), with the average weighted so that the sphere represents four additional points in the moment computation. The value of s is determined so that it is the average of the two components of the second moments at right angles to $X_o - X_s$, but limited by the average of all three components (the three eigenvalues), including the effect of 14 times the effect of the fictitious point, as above, as an upper limit, and excluding this effect, as a lower limit. This avoids putting undue weight on the $X_o - X_s$ dimension when the ellipsoid is long in this direction, since this dimension is less reliable because of the factor of 18 compression.

By combining the above information, the computation of the initial approximation can be expressed as follows:

$$\begin{aligned} X_s &= \frac{1}{n} \sum X_p \\ M_s &= \frac{1}{n} \sum (X_p - X_s)(X_p - X_s)^T \\ M_t &= 4M_s + \frac{14(X_o - X_s)(X_o - X_s)^T}{(X_o - X_s)^T M_s^{-1} (X_o - X_s)} \\ X_c &= X_s - \frac{2\sqrt{2}(X_o - X_s)}{\sqrt{(X_o - X_s)^T M_s^{-1} (X_o - X_s)}} \\ s &= \min \left[\max \left(\frac{4}{3} \text{tr}(M_s), \right. \right. \\ &\quad \left. \left. 2 \text{tr}(M_s) - 2 \frac{(X_o - X_s)^T M_s (X_o - X_s)}{(X_o - X_s)^T (X_o - X_s)}, \right. \right. \\ &\quad \left. \left. \frac{1}{3} \text{tr}(M_t) \right) \right] \\ M &= \frac{n}{n+4} M_t + \frac{4}{n+4} sI \\ W &= M^{-1} \end{aligned}$$

where X_p represents any point in the cluster, n is the number of points in the cluster, the summations are over these points, and I is the 3-by-3 identity matrix.

4. Iterative Solution for Ellipsoids

The adjustment of the ellipsoids is done by a modified least-squares approach. Each ellipsoid is adjusted so as to minimize the weighted sum of the squares of two kinds of discrepancies: the amounts by which the points (usually points in the cluster being fit) miss lying in surface of the ellipsoid, and the amounts by which the ellipsoid hides any points as seen from the camera position. (In the latter case, the discrepancies actually should be considered separately for each camera that sees the point in question. However, for narrow-angle stereo we use as a reasonable approximation the assumption that the "camera" is at the midpoint of the stereo baseline.) Including the second kind of discrepancy is useful in helping to determine the size and shape of the object when the points on the object itself do not contain sufficient information. Also included in the weighted sum of squares to be minimized are *a priori* terms which tend to force the ellipsoid by default to become a sphere near the ground when the points do not constrain it well.

The first kind of discrepancy above optimally should be defined as the length of the normal from the point in question to the surface of the ellipsoid. However, computing this requires solving a sixth-degree equation. Therefore, as an expedient the distance between the point and the surface along a straight line from the center of the ellipsoid to the point is used instead. In order to be consistent with this definition, the second kind of discrepancy is defined as follows. The midpoint of the two intersections of the surface of the ellipsoid with a line from the camera to the point is first found. Then the discrepancy of the first kind is computed for this midpoint. Both kinds of discrepancies are illustrated in Figures 1 and 2.

Now we must consider exactly for which points which kind of discrepancy is computed. There are five regions of space to consider, according to whether the point is to the side of the ellipsoid as seen from the camera (that is, the line through the camera position and the point does not intersect the ellipsoid), is in front of the ellipsoid as seen from the camera, is inside the front portion of the ellipsoid (in front of the surface of midpoints as defined above), is inside the back portion of the ellipsoid, or is behind the ellipsoid. Also, there are two kinds of points to consider, according to whether or not the point is in the cluster which is assumed to correspond to this object. This produces ten combinations in all, which are illustrated in Figures 1 and 2. They divide into four categories.

First, if the point is not in the cluster and is either in front of the ellipsoid or is to the side, there is no discrepancy and this point is not included in the computations.

Second, if the point is in the cluster and is either in front, inside the front half, or to the side, or if the point is not in the cluster and is inside the front half, the first kind of discrepancy is used.




Figure 1. Two-dimensional version of adjustment, showing points belonging to this cluster. Solid dark straight lines are discrepancies to be minimized.




Figure 2. Two-dimensional version of adjustment, showing points not in this cluster. Solid dark straight lines are discrepancies to be minimized.

Third, if the point is not in the cluster and is behind the ellipsoid, or if either kind of point is inside the back half, the second kind of discrepancy is used.

Fourth, if the point is in the cluster and is behind the the ellipsoid, both kinds of discrepancies are used, and the point acts as two points in the computations. This is because there are two separate components of error in this case: the object

does not extend enough on the side to hide the point, but it apparently bulges out in back (relative to the ellipsoid) to include the point.

In order to derive the mathematics for dividing space into the above five regions, consider the equation for the ellipsoid, as previously stated,

$$(X - X_c)^T W (X - X_c) = 1$$

and the equation of a straight line through the camera position X_o and the point in question X_p , in parametric form,

$$(X - X_o) = u(X_p - X_o)$$

These can be combined to produce

$$[u(X_p - X_o) + X_o - X_c]^T W [u(X_p - X_o) + X_o - X_c] = 1$$

the roots of which determine the intersections of the line and ellipsoid. This equation is equivalent to

$$au^2 + bu + c = 0$$

where

$$a = (X_p - X_o)^T W (X_p - X_o)$$

$$b = 2(X_p - X_o)^T W (X_o - X_c)$$

$$c = (X_o - X_c)^T W (X_o - X_c) - 1$$

The roots of the above equation in u determine the region of space in which X_p lies. If the roots are imaginary ($b^2 - 4ac < 0$), the point is to the side of the ellipsoid. If the average of the two roots ($-b/2a$) is positive, the point is in front of the midpoint surface, if negative, it is behind the surface. If the roots are real, the point is in front of, behind, or inside the ellipsoid according to whether both roots are greater than unity, both roots are less than unity, or unity lies between the roots, respectively. Alternatively, we can use the fact that the point is outside of the ellipsoid if and only if $(X_p - X_c)^T W (X_p - X_c) > 1$.

The discrepancy of the first kind is

$$\epsilon = \sqrt{(X_p - X_c)^T W (X_p - X_c)} \left(1 - \frac{1}{\sqrt{(X_p - X_c)^T W (X_p - X_c)}} \right)$$

In order to compute the discrepancy of the second kind, the midpoint of the intersections of the camera-point line with the ellipsoid is first obtained as follows:

$$X_u = -\frac{b}{2a}(X_p - X_o) + X_o$$

Then the discrepancy of the second kind is

$$\epsilon = \sqrt{(X_u - X_c)^T W (X_u - X_c)} \left(1 - \frac{1}{\sqrt{(X_u - X_c)^T W (X_u - X_c)}} \right)$$

Because there may be erroneous points in the data, points

which have large discrepancies relative to the size of the ellipsoid are given less weight in the solution. The weighting function used is

$$\omega = \frac{1}{\frac{1 + 2(\sqrt{(X-X_c)^T W (X-X_c)} - 1)^2}{\text{tr}(W)} + \sigma^2}$$

where X represents X_p or X_u for discrepancies of the first or second kinds, respectively, and σ is the component of standard deviation of measurement errors in X_p propagated into the discrepancy. (If these are unknown, σ can be zero.) Thus the dimensionless quantity to be minimized (by adjusting X_c and W) is $\Sigma \omega \epsilon^2$, plus some additional terms for *a priori* values yet to be discussed. However, this quantity is minimized only with respect to the effects of X_c and W acting through ϵ and not their effects through ω . In order to solve the above nonlinear problem, the Gauss method [6] is used. This method is equivalent to using the partial derivatives of the discrepancies to approximate the nonlinear problem by a linear statistical model [7], solving the linear problem, and iterating this process until it converges.

On any one iteration the following is done. The current values of X_c and W are used to compute for each point the value of ϵ as define above and the 1-by-9 matrix P , which consists of the partial derivatives of ϵ with respect to the three elements of X_c and the six unique elements of W . (W is symmetrical.) The following summations over all of the points are computed, in which each point in the first category above is not used, each point in the second or third categories appears once, and each point in the fourth category appears twice:

$$H = H_0 + \sum P^T \omega P$$

$$C = C_0 + \sum P^T \omega \epsilon$$

(H_0 and C_0 are used for the *a priori* values yet to be discussed.) Then the 9-by-1 matrix of corrections is

$$D = v H^{-1} C$$

where v is a factor used to improve convergence because of the very nonlinear nature of the problem. (Currently $v = 0.5$ on early iterations, but $v = 1$ after a test indicates that this will produce more rapid convergence.) The elements of D are subtracted from the corresponding elements of X_c and W to obtain the improved approximations for the next iteration.

Now the *a priori* values will be discussed. In some cases the points affecting the ellipsoid will be insufficient in number or insufficiently distributed to determine all parameters of the ellipsoid very well. It is therefore desirable to have *a priori* values for some of the parameters with appropriate weight in the solution to constrain them to reasonable default values when the points do not contain sufficient information. When there is ample information in the points, the *a priori* values will have very little effect because of their small weight. The *a priori* values currently used are the ground

surface height directly under X_c for the vertical component of X_c , with weight $0.1/\text{tr}(M)$, equality for the diagonal elements of W , with weight $\text{tr}(M)^2/10$, and zero for the off-diagonal elements of W , with weight $\text{tr}(M)^2/10$, where $M = W^{-1}$. (Including $\text{tr}(M)$ as shown scales things correctly so that the solution is invariant under a scale factor change.) The effect of the W terms is to try to force the ellipsoid into a spherical shape. These *a priori* terms are put into the solution in the following way. The diagonal element of H_0 corresponding to the vertical component of X_c is $0.1/\text{tr}(M)$, the three diagonal elements corresponding to the off-diagonal elements of W are each $\text{tr}(M)^2/10$, and the 3-by-3 submatrix on the diagonal of H_0 in the position corresponding to the diagonal elements of W consists of 2/3 on its main diagonal and -1/3 elsewhere multiplied by $\text{tr}(M)^2/10$. All other elements of the 9-by-9 matrix H_0 are zero. Then

$$C_0 = H_0 G$$

where G is a column matrix of the current values of X_c and W , arranged as in D , with the height of the ground directly under the center of the ellipsoid subtracted from the element of G corresponding to the vertical component of X_c . H_0 and C_0 are used in the summations for H and C as previously shown.

5. Breaking and Merging Clusters

Because the preliminary clustering is dependent on local information, it may not produce the best segmentation based on more global information. Therefore, after ellipsoids have been fit to all of the preliminary clusters, these clusters may be tentatively broken into smaller clusters and merged into larger clusters, new ellipsoids are fit to these clusters by the same process previously described, and a decision on whether to keep or reject each of these actions is made based on the goodness of fit of the ellipsoids to the points.

In order to decide where to break a cluster, for each edge in the portion of the original minimal spanning tree which connects this cluster the quantity $\lambda(1-\rho)$ is computed, where λ is the length of the edge and ρ is the minimum of

$\sqrt{(X_p - X_c)^T W (X_p - X_c)}$ for the two points connected by the edge. Then the cluster is tentatively broken at the edge for which this quantity is maximum, of all such edges such that each new cluster formed has at least four points at least one of which has $(X_p - X_c)^T W (X_p - X_c) > 1$ (that is, it is outside the old ellipsoid). This process tends to break the cluster at places furthest inside the ellipsoid, but connecting points that are outside the ellipsoid. If this new clustering is accepted by the criteria described below, the process repeats on the new clusters.

After the above breaking process is finished, any two clusters are tentatively merged if $(X_a - X_c)^T W (X_a - X_c) < 4$ for either cluster, where X_a is X_c for the other cluster, provided that these two clusters were not previously one cluster before breaking. If there is competition for the merging, the cluster

pair with the minimum value for this quantity is merged first. If a merger is accepted, further mergers can take place on these clusters by this same process.

The criteria for accepting two clusters or one that resulted from a tentative break or merger are as follows. If $(X_a - X_c)^T W (X_a - X_c) < 1$ for either small cluster, where X_a is X_c for the other small cluster (that is, the center of one ellipsoid is inside the other ellipsoid), the single cluster is chosen. Otherwise, the following quantity is computed for each of the three ellipsoids:

$$q = \frac{\sum \omega \epsilon^2}{n-3}$$

where ϵ and ω are the discrepancies and weights from the last iteration, as defined in the previous section, and n is the number of points in the cluster corresponding to this ellipsoid. (If the initial approximation is used as the result, ϵ and ω are obtained from the first iteration, and the denominator is n instead of $n-3$.) Then the two small clusters are chosen if the sum of their two values of q is less than the value of q for the single cluster. Otherwise, the single cluster is chosen.

6. Results

Fig. 3 shows a stereo pair of pictures taken from the Viking Lander 1 on the surface of Mars. Each picture is 256 pixels by 256 pixels. The pixel spacing is 0.04 degrees in azimuth and elevation. (Thus the field of view is about 10 degrees.) The azimuth and elevation from the left camera to the center of the picture are about 18 degrees and -20 degrees, respectively, relative to the camera positions. The two cameras are 0.8187 meters apart. The distances to the points in the scene range from about 3 meters to about 4.5 meters.

Fig. 4 shows the left picture enlarged, with arrows indicating the points found by the stereo processing. The point of each arrow is at the center of an eight-pixel-square window which was matched to a corresponding area in the other picture. (Thus the resolution of the reduced stereo data is 0.32 degrees in azimuth and elevation.) The arrows are dropped perpendicularly from the point in three-dimensional space computed for this point to a nominally horizontal reference plane 1.5 meters below the cameras, with the base of the arrows on this plane, and are then projected into the picture. Fig. 5 shows the same data as Fig. 4 except that the bases of the arrows rest on a ground plane computed from this data by the ground surface finder. Fig. 6 is the same as Fig. 5 except that it shows only points computed to be at least 5 centimeters above the ground surface.

The 5-centimeter height threshold was used for selecting the points to cluster in the object finder. The minimum distance for breaking the minimal spanning tree to form the initial clusters was also 5 centimeters, and the maximum distance for connecting points was 20 centimeters. (Using zero and infinity for this minimum and maximum produced a slightly different initial clustering but identical final results.)

Fig. 7 shows the points in Fig. 6 in a nominally vertical orthogonal projection (perpendicular to the reference plane). The figure covers an area one meter by one meter. The symbol for each point represents height in centimeters above the ground plane, with the letters A, B, C, etc. representing the values 10, 11, 12, etc. The points are connected to show the minimal spanning trees that were computed. Solid lines connect points within each initial cluster.




Figure 3. Stereo pair of Martian surface.




Figure 4. Points found by stereo processing, showing heights above reference plane.




Figure 5. Heights above computed ground plane.




Figure 6. Heights above computed ground plane, for points above 5 cm.




Figure 7. Minimal spanning tree connecting points above 5 cm, vertical view. Solid lines show initial clusters.




Figure 8. Ellipsoids fit to initial clusters.




Figure 9. Ellipsoids fit to final clusters.




Figure 10. Ellipsoids fit to final clusters, projected into left picture.

Fig. 8 shows the ellipsoids that were fit to the initial clusters. Each ellipsoid is represented by two ellipses. One ellipse is the orthogonal projection of the ellipsoid onto the reference plane. The other ellipse is the intersection of the ellipsoid with a plane through the center of the ellipsoid and parallel to the reference plane. Only the clustered points are shown here, as in Fig. 7. However, as previously described, any of the points shown in Fig. 4 may have been involved in the adjustment of the ellipsoids. Remember that the fit is done in three dimensions, whereas Fig. 8 shows a two-dimensional projection.

Fig. 9 shows in the same way the results of the breaking and merging operations. The two clusters in the center (corresponding to the large rock in the center of the pictures) were merged into one, and a new ellipsoid is shown for this cluster. The other clusters were not changed.

These results were projected into the left picture to produce Fig. 10. The outline of the ellipsoids as they would be seen from the left camera are superimposed on the picture. The lengths of the principal axes of the large ellipsoid in the center are 36.5, 30.5, and 19.8 centimeters.

Acknowledgment

The Mars pictures were supplied by Elliott Levinthal and Sidney Liebes of the Viking Lander Imaging Team.

References

- [1] Gennery, D.B. "A Stereo Vision System for an Autonomous Vehicle." In *Proc. IJCAI-77*. MIT, Cambridge, Mass., August 1977, pp. 576-582.
- [2] Nevatia, R. & Binford, T.O. "Description and Recognition of Curved Objects," *Artif. Intell.* Vol. 8, February 1977, pp. 77-98.
- [3] Hohn, F.E. *Elementary Matrix Algebra* (Third Edition). The MacMillan Company, 1973.
- [4] Zucker, S.W. "Relaxation Labelling and the Reduction of Local Ambiguities," In *Proc. Third International Joint Conf. on Pattern Recognition*. San Diego, Calif., November 1976, pp. 852-861.
- [5] Duda, R. & Hart, P. *Pattern Recognition and Scene Analysis*. Wiley, 1973.
- [6] Bard, Y. *Nonlinear Parameter Estimation*. Academic Press, 1974.
- [7] Graybill, F.A. *An Introduction to Linear Statistical Models*, Volume I. McGraw-Hill Book Company, 1961.