

AD-A102 222

ROCHESTER UNIV NY DEPT OF COMPUTER SCIENCE
EXPERIENCE WITH THE GENERALIZED HOUGH TRANSFORM, (U)
OCT 80 K R SLOAN, D H BALLARD
TR-83

F/G 9/2

N00014-78-C-0164

NL

UNCLASSIFIED

1 1 1
AD
A102222



END
DATE
FILMED
8 81
DTIC

AD A102222

LEVEL II

13

DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

Rochester

Department of Computer Science
University of Rochester
Rochester, New York 14627

DTIC
SELECTED
JAN 30 1981

THIS FILE COPY

81 6 15 17

12 13

4171R-83

EXPERIENCE WITH THE GENERALIZED
HOUGH TRANSFORM

(10) K.R. Sloan, Jr.
D.H. Ballard
Computer Science Department
The University of Rochester
Rochester, NY 14627

(1) TR-83
October 1980

Accession For	
NTIS GRA&I	☒
DTIC TAB	☐
Unannounced	☐
Justification	
Per Little on file	
By 1422	
Distribution/	
Availability Codes	
Avail and/or	
Dist	Special
A	

Abstract

The Hough Transform is a method for detecting curves by exploiting the duality between points on a curve and parameters of that curve. The initial work showed how to detect both analytic curves [Hough, 1962; Duda and Hart, 1972] and non-analytic curves [Merlin and Farber, 1975], in the case of binary edge images. This work was generalized to the detection of some analytic curves in grey level images, specifically lines [O'Gorman and Clowes, 1973], circles [Kimme et al., 1975], and parabolas [Wechsler and Sklansky, 1977].

Recently, the Hough technique has been extended to the detection of arbitrary non-analytic shapes in grey level images [Ballard, 1979]. This shape detection scheme has been implemented and tested on a variety of artificial images and has found application in the analysis of real aerial images. Experience to date indicates that the technique is robust with respect to occlusions, but requires reliable edge-element orientation determination.

15

The preparation of this paper was supported in part by the Defense Advanced Research Projects Agency, monitored by the ONR, under Contract N00014-78-C-0164.

410370

1. Introduction

Shape is an important attribute of two-dimensional figures. In simple figure-ground binary images, the shape of the boundary of the figure is often the only interesting feature. We take "shape" to be a property of the entire figure, i.e., it is a global property.

Evidence about the shape of a figure is found at the boundary between figure and ground. Such evidence can be generated by the application of local edge-element detectors. An edge-element detector typically reports on the presence of an edge-element in a small window of an image, and on the orientation of that edge-element. Finding shapes in the image involves combining many pieces of local evidence into a global judgment.

The Hough Transform is a method for detecting curves by exploiting the duality between points on a curve and parameters of that curve. The initial work showed how to detect both analytic curves [Hough, 1962; Duda and Hart, 1972] and non-analytic curves [Merlin and Farber, 1975], in the case of binary edge images. This work was generalized to the detection of some analytic curves in grey level images, specifically lines [O'Gorman and Clowes, 1973], circles [Kimme et al., 1975], and parabolas [Wechsler and Sklansky, 1977].

Recently, the Hough technique has been extended to the detection of arbitrary non-analytic shapes in grey level images [Ballard, 1979]. Given an arbitrary shape, S , this generalized Hough technique provides a mapping from the orientation of an edge-element to the set of instances of S (as modified by location, rotation, and uniform scaling) which could have given rise to that edge-element. This mapping allows all local evidence for a particular instance of S to contribute to global decisions about the figure.

This shape detection scheme has been implemented and tested on a variety of artificial images and has found application in the analysis of real aerial images. Experience to date indicates that the technique is robust. Also, with appropriate "focus of attention" mechanisms, which are present in our implementation, the method is also efficient. However, the reliable determination of edge-element orientation is crucial to the success of this method.

2. Hough Techniques

All Hough techniques for shape detection consist of the following basic elements:

- a) a local edge-element detector, E ,
- b) an n -dimensional parameter space, P , quantized and represented by an n -dimensional

- Accumulator Array, AA,
- c) a mapping, M, from the information provided by E into P (and thus AA),
 - d) a voting rule, V, specifying how a particular edge-element affects the values of AA,
 - e) a Detection rule, D, specifying the conditions under which a particular shape has been detected.

Given these basic elements, shapes are found by the following procedure:

- a) zero AA,
- b) apply E everywhere in the image,
- c) for each edge-element found, apply M to locate cells in AA. Then apply V to modify the contents of these cells. (i.e., vote for all possible "causes" of this edge-element),
- d) finally, apply D to AA (choose the most popular shape).

Clearly, application of this technique depends on the ability to parameterize the shapes of interest, and the derivation of the mapping M from edge-element information to possible shape parameters.

Lines

The original Hough transform capitalized on the observation that straight lines can be completely specified by two parameters (e.g., an orientation [theta], and a distance from the origin, s). What is more, the mapping, from a particular edge-element position to the set of straight lines it might be a part of, is easy to compute [Hough, 1962; Duda and Hart, 1972]. The idea is that an actual line in the image will give rise to many local edge-elements, all of which will "vote" for that line. Individual edge-elements will also vote for other lines, but the "correct" line will receive the most votes.

If the edge-element operator, E, provides directional information, then each edge-element maps to a unique line. Edge elements which line up vote for "their" line, and the line with the most visible edge-elements gets the most votes. Note that it is not necessary for the edge-elements to be connected (or even be near each other) in order that their votes reinforce one another--they must simply be colinear.

Circles

The description of circular figures in an image requires three parameters: x, y, s. The location of the center of the circle is given by <x,y> and the radius is given by the scale parameter, s. Each edge-element in the image is evidence for a set of <x,y,s> triples.

If the direction of the edge-element is unknown, then the locus of points in parameter space representing circles which could have created this edge-element forms a right circular cone. In the presence of direction information, this locus is reduced to a line [Ballard, 1979]. As with line detection, circles which actually appear in the image will receive many votes; those which do not will receive few votes.

Arbitrary Shapes

The Hough technique can be extended to analytic shapes for which the mapping from edge-element to a locus of points in parameter space can be derived. Given certain assumptions about the meaning of "shape," we can also extend the technique to arbitrary, non-analytic shapes.

Consider a particular figure (e.g., an ellipse centered at $\langle 1, 2 \rangle$ with its major axis parallel to the x-axis and of length 10, and its minor axis of length 5). Now, consider the set of figures which can be produced by translating, rotating, and uniformly scaling the original figure. For our purposes, all of these figures have the same shape.

The parameter space which captures this notion of shape is:

$$P = \langle x, y, s, [\text{theta}] \rangle$$

where $\langle x, y \rangle$ is the origin of a local co-ordinate system, s is a scale factor, and $[\text{theta}]$ is a rotation about $\langle x, y \rangle$. This is the parameter space used in our generalized Hough Transform. Note that the Hough-spaces developed above for lines and circles are sub-spaces of P .

The key to all Hough techniques is the mapping from edge-element information to a locus of points in P . We assume an edge-element operator which provides directional information. As seen above, this directional information can drastically reduce the image of the edge-element in P . Our mapping, M , depends strongly on the reliability of the edge-element direction.

Consider the hyperplane of P with

$$[\text{theta}] = 0, s = 1.$$

We represent the mapping from edge-element location and orientation to figure location directly in an "R-Table" (see Figures 1 & 2). The orientation of an edge-element is used as an index into this table, where are stored a set of $\langle x, y \rangle$ vectors. When added to the $\langle x, y \rangle$ location of the edge-element in the image, these vectors point to possible locations for the origin of a figure's local co-ordinate system (its reference point). This map is easy to build, given an original master shape.

The expansion of the R-Table mapping to cover the remainder of P is performed dynamically by our voting procedure, V. This involves rotating the edge-element orientation before using it as an index into the R-Table, and scaling the R-Table entries thus formed before calculating the figure's hypothesized reference point.

3. Implementation and Experimental Results

The generalized Hough Transform described above has been implemented and tested on a variety of artificial images and has found application in the analysis of real aerial images. Experience to date indicates that the technique is robust, given that the edge-element operator used to generate local evidence for the shape can provide reliable information about edge-element direction.

R-Tables

The R-Table defines the mapping from edge-element information (position and orientation) into a hyperplane of parameter space. This mapping is derived from an explicit master shape, in the form of a sequence of boundary points. Typically, we sketch (or trace) a shape. In order to ease the pain of carefully drawing a particular shape, we customarily sample the master shape boundary rather coarsely and then fill it with a B-spline fit to these points [Riesenfeld, 1973]. An arbitrary reference point is chosen for the origin of the local co-ordinate system.

Now, for each point on the master shape boundary, we calculate the orientation of the boundary edge-element at that point and the vector from the boundary edge-element to the origin of the local co-ordinate system. This is exactly an R-Table entry. The current implementation of the R-Table consists of a list of entries, tagged with the edge-element orientation, containing a list of reference point vectors. Any scheme which associates edge-element orientation with reference point vectors will do.

Edge Detection

The examples shown below used a simple 3x3 Sobel edge-element finder. In general, this is satisfactory. When, as in one example below, this does not provide reliable edge-element orientation, performance deteriorates seriously.

Detection Criteria

For the purposes of these examples, the shape found by the generalized Hough Transform is determined by simply selecting the maximum value found in a smoothed (over a 3x3x3x3 window) Accumulator Array. This does the right thing when, as in most of our examples, the maxima in the Accumulator Array are sharp peaks. For more problematic, noisy situations, clustering in

parameter space may be required.

Artificial Images

Figure 3a-3d illustrates a few of the features of the experimental implementation of the generalized Hough Transform. These artificial images provide controlled conditions for our testing. In Figure 3a and 3b we see that, as expected, the method has no difficulty in finding the central shape at arbitrary scale and orientation (The black dots show the shape, as drawn from the R-Table and the parameter choices which received the most votes in the Accumulator Array, the central black dot is the reference point.) In Figure 3c we see what appears to be the same shape, obscured by another. Figure 3d demonstrates that there is enough evidence for the desired shape to correctly determine its location, orientation, and scale.

All of these images, of course, have very clean edges and the 3x3 Sobel operator has no difficulty in correctly determining edge-element orientation. By way of contrast, see Figure 4a. In this image, which has been degraded by the addition of Gaussian noise of mean zero and standard deviation ten, the edge-elements found by the 3x3 Sobel operator are too short, and the noise hopelessly jumbles the orientation information. As a result, the generalized Hough Transform (which depends strongly on the accuracy of edge-element orientation) is unable to locate the shape. The guess shown is not much more than that--the Accumulator Array has no very strong peak, and we simply show the shape instance which received the most votes in a very close election.

To ameliorate the effects of noise, the image can be smoothed prior to applying the edge operator. As an experiment, the noisy image of Fig. 4a was smoothed by convolving it with a 5x5 template of ones. Next, the Hough algorithm was applied as before. Fig. 4b shows that, in this case, the edge estimates have been improved enough so that the shape is now correctly located.

Aerial Photographs

The location of arbitrary, non-analytic shapes is not merely of interest in artificial images such as that shown above. The original version of the shape found above came from the aerial image shown as Figure 5a. Even the experimental version of the generalized Hough Transform has no difficulty in locating the pond in this image, as shown in Fig. 5b.

4. Focus of Attention

One of the difficulties encountered in the application of this technique to real images, for the location of real shapes, is that the area searched for evidence of boundaries (the application of the edge-element detector and the mapping from

edge-element information to parameter space) and the size of the parameter space can quickly become very large. The solution to these problems is, of course, to attempt to focus attention where possible.

Where to Look

One way to focus attention during the application of this shape-finding technique is to constrain the area searched for evidence of the figure's boundary [Russell and Brown, 1978; Russell, 1979]. In a system which routinely applies an edge operator over the entire image, this may not seem to be a solution (or even a problem). However, even after the edge-elements have been found, it is still necessary to apply the mapping to parameter space (one per desired shape). Our implementation includes the usual "bounding rectangle" limitation on the area in which edge-elements are to be found and mapped to parameter space. This improves performance significantly.

What to Look For

The second obvious way to focus attention is to constrain the objects being sought. Of course, a single application of the generalized Hough Transform concentrates on the location of a particular class of shape (that defined by the R-Table). In addition, it is usually possible to constrain the permissible values for some (if not all) of the parameters. Constraining the location of the reference point is related to the question of "Where to look." Constraining the parameters of scale or rotation is also possible, and certainly worth doing. Sometimes, the unconstrained search for a particular shape (such as the pond in Figure 5) will result in almost complete information about the range of values to be considered in successive searches. For example, once the pond has been located (in parameter space, including location in the image, rotation and scale) map-like knowledge about this particular part of the world would allow the search for other shapes in the scene to be almost completely determined. Thus, our technique can both generate and benefit from such constraints.

Although our current implementation uses only a simple "bounding rectangle" constraint on the area of the image to be searched for boundary information, it is possible to combine information about the range of locations for the reference point, scale, and rotation. When all of these are sufficiently constrained, then the R-Table itself provides pointers to the locations to be searched in the image for edge-elements.

5. Conclusion

Shape is an important defining feature of many image objects, often the only useful feature. The key ideas behind the Hough Transform have been extended to produce a shape detection technique which performs well in the presence of occlusion, even

for completely arbitrary, non-analytic shapes. As has been demonstrated, however, the technique depends strongly on the reliable estimation of edge-element orientation.




Figure 1: Geometry for Generalized Hough Transform.

i	ϕ_i	
0	0	$\{r \mid a - r = x, x \text{ in } B, \phi(x) = 0\}$
1	$\Delta\phi$	$\{r \mid a - r = x, x \text{ in } B, \phi(x) = \Delta\phi\}$
2	$2\Delta\phi$	$\{r \mid a - r = x, x \text{ in } B, \phi(x) = 2\Delta\phi\}$
...	...	...

Figure 2: R-Table Format.





Figure 3a, 3b





Figure 3c, 3d





Figure 4a, 4b



Figure 5a, 5b

REFERENCES

- Ballard, D.H. and Sklansky, J., A ladder-structured decision tree for recognizing tumors in chest radiographs, IEEE Transactions on Computers 25, 1976, 503-513.
- Ballard, D.H., Generalizing the Hough transform to detect arbitrary shapes, TR55, Computer Science Department, University of Rochester, October 1979; submitted to Pattern Recognition.
- Duda, R.O. and Hart, P.E., Use of the Hough transform to detect lines and curves in pictures, CACM 15, 1, January 1972, 11-15.
- Hough, P.V.C., Method and means for recognizing complex patterns, U.S. Patent 3,069,654, 1962.
- Kimme, C., Ballard, D.H., and Sklansky, J., Finding circles by an array of accumulators, CACM 18, 1, February 1975, 120-122.
- Merlin, P.M. and Farber, D.J., A parallel mechanism for detecting curves in pictures, IEEE Transactions on Computers 24, January 1975, 96-98.
- O'Gorman, F., and Clowes, M.B., Finding picture edges through collinearity of feature points, Proc., Third IJCAI, Stanford, CA, August 1973, 543-555.
- Riesenfeld, R., Application of B-Spline approximation to geometric problems of Computer-Aided-Design, UTEC-CSs-73-126, University of Utah, March 1973.
- Russell, D.M., Where do I look now?: Modeling and inferring object locations by constraints, Proc., IEEE Conference on Pattern Recognition and Image Processing, Chicago, IL, August 1979.
- Russell, D.M. and Brown, C.M., Representing and using locational constraints in aerial imagery, DARPA Image Understanding Workshop, November 1978.
- Sklansky, J., On the Hough technique for curve detection, IEEE Transactions on Computers, July 1977.
- Shapiro, S.D., Transformation for the computer detection of curves in noisy pictures, Computer Graphics and Image Processing 4, 1975, 328-338.
- Shapiro, S.D., Properties of transforms for the detection of curves in noisy pictures, Computer Graphics and Image Processing 8, 1978, 219-236.

Wechsler, H. and Sklansky, J., Automatic detection of ribs in chest radiographs, Pattern Recognition 9, January 1977, 21-30.

END

DATE
FILMED

8 - 81

DTIC