

GENERALIZING THE HOUGH TRANSFORM TO DETECT ARBITRARY SHAPES*

D. H. BALLARD

Computer Science Department, University of Rochester, Rochester, NY 14627, U.S.A.

(Received 10 October 1979; in revised form 9 September 1980; received for publication 23 September 1980)

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^(1,2) and non-analytic curves,⁽³⁾ but these methods were restricted to binary edge images. This work was generalized to the detection of some analytic curves in grey level images, specifically lines,⁽⁴⁾ circles⁽⁵⁾ and parabolas.⁽⁶⁾ The line detection case is the best known of these and has been ingeniously exploited in several applications.^(7,8,9)

We show how the boundaries of an *arbitrary* non-analytic shape can be used to construct a mapping between image space and Hough transform space. Such a mapping can be exploited to detect instances of that particular shape in an image. Furthermore, variations in the shape such as rotations, scale changes or figure-ground reversals correspond to straightforward transformations of this mapping. However, the most remarkable property is that such mappings can be composed to build mappings for complex shapes from the mappings of simpler component shapes. This makes the generalized Hough transform a kind of universal transform which can be used to find arbitrarily complex shapes.

Image processing
Parallel algorithms

Hough transform

Shape recognition

Pattern recognition

1. INTRODUCTION

In an image, the pertinent information about an object is very often contained in the shape of its boundary. Some appreciation of the importance of these boundary shapes in human vision can be gained from experiments performed on the human visual system, which have shown that crude encodings of the boundaries are often sufficient for object recognition⁽¹⁰⁾ and that the image may be initially encoded as an 'edge image', i.e. an image of local intensity or color gradients. Marr⁽¹¹⁾ has termed this edge image a 'primal sketch' and suggested that this may be a necessary first step in image processing. We describe a very general algorithm for detecting objects of a specified shape from an image that has been transformed into such an edge representation. In that representation, sample points in the image no longer contain grey level information, but instead each sample point contains a magnitude and direction representing the severity and orientation of the local grey level change.

Operators that transform the image in such a way are known as edge operators, and many such operators are available, all based on different models of the local grey level changes. Two of the most used are the gradient operator (for example, see Prewitt⁽¹²⁾) and the Hueckel operator,⁽¹³⁾ which model local grey level changes as a ramp and a step respectively.

Our generalized Hough algorithm uses edge information to define a mapping from the orientation of an edge point to a reference point of the shape. The reference point may be thought of as the origin of a local co-ordinate system for the shape. Then there is an easy way of computing a measure which rates how well points in the image are likely to be origins of the specified shape. Figure 1 shows a few graphic examples of the information used by the generalized Hough transform. Lines indicate gradient directions. A feature of the transform is that it will work even when the boundary is disconnected due to noise or occlusions. This is generally not true for other strategies which track edge segments.

The original algorithm by Hough⁽²⁾ did not use




Fig. 1. Kinds of shapes detected with generalized Hough transform. (a) Simple shape; (b) composite shape.

* The research described in this report was supported in part by NIH Grant R23-HL-2153-01 and in part by the Alfred P. Sloan Foundation Grant 78-4-15.

orientation information of the edge, and was considerably inferior to later work using the edge orientation for parametric curves.^(5,6,14) Shapiro^(15,16,17) has collected a good bibliography of previous work as well as having contributed to the error analysis of the technique.

1.1 Organization

Section 2 describes the Hough transform for analytic curves. As an example of the parametric version of the transform, we use the ellipse. This example is very important due to the pervasiveness of circles in images, and the fact that a circle becomes an ellipse when rotated about an axis perpendicular to the viewing angle. Despite the importance of ellipses, not much work has used the Hough transform. The elliptical transform is discussed in detail in Section 3. Section 4 describes the generalized algorithm and its properties. Section 5 describes special strategies for implementing the algorithm and Section 6 summarizes its advantages.

2. THE HOUGH TRANSFORM FOR ANALYTIC CURVES

We consider analytic curves of the form $f(\mathbf{x}, \mathbf{a}) = 0$ where \mathbf{x} is an image point and \mathbf{a} is a parameter vector.




Fig. 2. (a) Locus of parameters with no directional information. (b) Locus of parameters with directional information.

To see how the Hough transform works for such curves, let us suppose we are interested in detecting circular boundaries in an image. In Cartesian coordinates, the equation for a circle is given by

$$(x-a)^2 + (y-b)^2 = r^2. \quad (1)$$

Suppose also that the image has been transformed into an edge representation so that only the magnitude of local intensity changes is known. Pixels whose magnitude exceeds some threshold are termed *edge pixels*. For each edge pixel, we can ask the question: if this pixel is to lie on a circle, what is the locus for the parameters of that circle? The answer is a right circular cone, as shown in Fig. 2(a). This can be seen from equation (1) by treating x and y as fixed and letting, a , b , and r vary.

The interesting result about this locus in parameter space is the following. If a set of edge pixels in an image are arranged on a circle with parameters a_0 , b_0 , and r_0 , the resultant loci of parameters for each such point will pass through the same point (a_0, b_0, r_0) in parameter space. Thus many such right circular cones will intersect at a common point.

2.1 Directional information

We see immediately that if we also use the *directional* information associated with the edge, this reduces the parameter locus to a line, as shown in Fig. 2(b). This is because the center of the circle for the point (x, y) must lie r units along the direction of the gradient. Formally, the circle involves 3 parameters. By using the equation for the circle together with its derivative, the number of free parameters is reduced to one. Formally, what happens is the equation

$$\frac{df}{dx}(\mathbf{x}, \mathbf{a}) = 0$$

introduces a term dy/dx which is known since

$$\frac{dy}{dx} = \tan \left[\phi(\mathbf{x}) - \frac{\pi}{2} \right]$$

where $\phi(\mathbf{x})$ is the gradient direction. This suggests the following algorithm.

Hough algorithm for analytic curves in grey level images. For a specific curve $f(\mathbf{x}, \mathbf{a}) = 0$ with parameter vector \mathbf{a} , form an array $A(\mathbf{a})$, initially set to zero. This array is termed an accumulator array. Then for each edge pixel \mathbf{x} , compute all \mathbf{a} such that $f(\mathbf{x}, \mathbf{a}) = 0$ and $df/dx(\mathbf{x}, \mathbf{a}) = 0$ and increment the corresponding accumulator array entries:

$$A(\mathbf{a}) := A(\mathbf{a}) + 1.$$

After each edge pixel \mathbf{x} has been considered, local maxima in the array A correspond to curves of f in the image.

If only the equation $f(\mathbf{x}, \mathbf{a}) = 0$ is used, the cost of the computation is exponential in the number of parameters minus one, that is, where m parameters each have M values, the computation is proportional to




Fig. 3. Using convolution templates to compensate for errors.

M^{m-1} . This is because the equation of the curve can be used to determine the last parameter. The use of gradient directional information saves the cost of another parameter making the total effort proportional to M^{m-2} , for $m \geq 2$.

2.2 Compensating for errors

A problem arises in detecting maxima in the array $A(\mathbf{a})$. Many sources of error effect the computation of the parameter vector \mathbf{a} so that in general many array locations in the vicinity of the ideal point \mathbf{a} are incremented instead of the point itself. One way of handling this problem is to use a formal error model on the incrementation step. This model would specify a set of nearby points instead of a single point. Shapiro⁽¹⁵⁻¹⁸⁾ has done extensive work on this subject. Another solution to this problem is to replace uncompensated accumulator values by a function of the values themselves and nearby points after the incrementation step. The effect of this operation is to smooth the accumulator array. We show that, under the assumption of isotropic errors, these methods are equivalent.

Returning to the initial example of detecting circles, the smoothing of the accumulator array is almost equivalent to the change in the incrementing procedure we would use to allow for uncertainties in the gradient direction ϕ and the radius r . If we recognized these uncertainties as:

$$\begin{aligned} \phi(\mathbf{x}) &\pm \Delta\phi \\ r &\pm \Delta r(r) \end{aligned}$$

we would increment all values of \mathbf{a} which fall within the shaded band of Fig. 3. We let Δr increase with r so that uncertainties are counted on a percentage basis. Figure 3 shows the two-dimensional analog of the general three-dimensional case.

Suppose we approximate this procedure by incrementing all values of \mathbf{a} which fall inside the square domain centered about the nominal center shown in Fig. 3, according to some point spread function h . After the first contributing pixel which increments center \mathbf{a}_0 has been taken into account, the new accumulator

array contents A will be given by

$$A(\mathbf{a}) = h(\mathbf{a} - \mathbf{a}_0) \quad (2)$$

where $\mathbf{a} = (a_1, a_2, r)$ and $\mathbf{a}_0 = (a_{10}, a_{20}, r_0)$. If we include all the contributing pixels for that center, denoted by C , the accumulator is

$$A(\mathbf{a}) = C(\mathbf{a}_0)h(\mathbf{a} - \mathbf{a}_0). \quad (3)$$

Finally for all incremented centers, we sum over \mathbf{a}_0 :

$$A(\mathbf{a}) = \sum_{\mathbf{a}_0} C(\mathbf{a}_0)h(\mathbf{a} - \mathbf{a}_0). \quad (4)$$

But $C(\mathbf{a}_0) = A(\mathbf{a}_0)$, so that

$$\begin{aligned} A(\mathbf{a}) &= \sum_{\mathbf{a}_0} A(\mathbf{a}_0)h(\mathbf{a} - \mathbf{a}_0) \\ &= A * h \\ &\equiv A_s(\mathbf{a}). \end{aligned} \quad (5)$$

Thus within the approximation of letting the square represent the shaded band shown in Fig. 3, the smoothing procedure is equivalent to an accommodation for uncertainties in the gradient direction and radius.

3. AN EXAMPLE: ELLIPSES

The description of the algorithm in Section 2.1 is very terse and its implementation often requires considerable algebraic manipulation. We use the example of finding ellipses to show the kinds of calculation which must be done. Ellipses are an important example, as circles, which are a ubiquitous part of many everyday objects, appear as ellipses when viewed from a distant, oblique angle.

We use the center of the ellipse as a reference point and assume that it is centered at x_0, y_0 with major and minor diameters a and b . For the moment, we will assume that the ellipse is oriented with its major axis parallel to the x -axis. Later we will relax this requirement by introducing an additional parameter for arbitrary orientations. For the moment, assume a and b are fixed. Then the equation of the ellipse is:

$$\frac{(x-x_0)^2}{a^2} + \frac{(y-y_0)^2}{b^2} = 1. \quad (6)$$




Fig. 4. Parametrization of an ellipse with major axis parallel to x -axis.

Let $X = x - x_0$, $Y = y - y_0$, then

$$\frac{X^2}{a^2} + \frac{Y^2}{b^2} = 1 \quad (7)$$

Differentiating with respect to X

$$\frac{2X}{a^2} + \frac{2Y}{b^2} \frac{dY}{dX} = 0. \quad (8)$$

But dY/dX is known from the edge pixel information! Let $dY/dX = \xi$, then from (8)

$$X^2 = \left(\frac{a^2}{b^2} \xi\right)^2 Y^2. \quad (9)$$

Substituting in (7)

$$\frac{Y^2}{b^2} \left(1 + \frac{a^2}{b^2} \xi^2\right) = 1 \quad (10)$$

$$Y = \pm \frac{b^2}{\sqrt{\left(1 + \frac{a^2}{b^2} \xi^2\right)}} \quad (11)$$

so that

$$X = \pm \frac{a^2}{\sqrt{\left(1 + \frac{b^2}{a^2} \xi^2\right)}} \quad (12)$$

and finally, given a, b, x, y and dY/dX , we can determine x_0 and y_0 as:

$$x_0 = x \pm \frac{a^2}{\sqrt{\left(1 + \frac{b^2}{a^2} \xi^2\right)}} \quad (13)$$

$$y_0 = y \pm \frac{b^2}{\sqrt{\left(1 + \frac{a^2}{b^2} \xi^2\right)}} \quad (14)$$

The four solutions correspond to the four quadrants, as shown in Fig. 5. The appropriate quadrant can be found from the gradient by testing the signed differences dY and dX .

The final step is to handle rotations by introducing a fifth parameter θ . For an arbitrary θ , we calculate (X, Y) using

$$\xi = \tan\left(\phi - \theta - \frac{\pi}{2}\right)$$

and rotate these (X, Y) by θ to obtain the correct




Fig. 5. Four reference point solutions resolvable with gra-

(x_0, y_0) . In ALGOL we would implement this as:

```

procedure HoughEllipse (integer Xmin, Xmax, Ymin, Ymax,
θmin, θmax, amin, amax, bmin, bmax, x, y, x0, y0, dx, dy; real angle, ξ;
integer array A, P);
begin
for x := xmin step dx to xmax do
for y := ymin step dy to ymax do
begin
dX := P(x + delta, y) - P(x, y);
dY := P(x, y + delta) - P(x, y);
for a := amin step da until amax do
for b := bmin step db until bmax do
for θ := θmin step dθ until θmax do
begin
angle := arctan(dY/dX) - θ - π/2;
ξ := tan(angle);
dx := Sign X(dx, dy) * a^2 / sqrt(1 + b^2/ξ^2);
dy := Sign Y(dx, dy) * b^2 / sqrt(1 + a^2/b^2);
Rotate-by-Theta(dx, dy);
x0 := x + dx;
y0 := y + dy;
A(x0, y0, θ, a, b) := A(x0, y0, θ, a, b) + 1;
end;
end.

```

Notice that to determine the appropriate formulae for an arbitrary orientation angle θ , we need only rotate the gradient angle and the offsets dx and dy . $\text{Sign } X$ and $\text{Sign } Y$ are functions which return ± 1 depending on the quadrant determined by dX and dY .

3.1 Parameter space image space trade-offs

Tsuji and Matsumoto⁽¹⁹⁾ recognized that a decreased computational effort in parameter space could be traded for an increased effort in edge space. It is our intent to place these ideas on a formal footing. Later we will see that the same kind of trade-off is potentially available for the case of arbitrary shapes, but is impractical to implement.

An ellipse has five parameters. Referring to the basic algorithm in Section 2.1, we use the equation for the ellipse together with its derivative to solve for two of these parameters as a function of the other three. Thus the algorithm examines every edge point and uses a three-dimensional accumulator array so that the computations are of order $O(ed^3)$. Here e is the number of edge pixels and we are assuming d distinct values for each parameter. Suppose we use pairs of edge points in the algorithm. This results in four equations, two involving the equation for an ellipse evaluated at the different points and two for the related derivatives. This leaves one free parameter. Thus the resultant computational effort is now $O(e^2d)$. The detailed derivation of this form of the Hough algorithm is presented in the Appendix.

If parameter space can be highly constrained so that the set of plausible values is small, then the former

Table 1. Analytic curves described in terms of the generalized shape parameters $x_r, y_r, S_x, S_y, \theta$

Analytic form	Parameters	Equation
Line	S, θ	$x \cos \theta + y \sin \theta = S$
Circle	x_r, y_r, S	$(x - x_r)^2 + (y - y_r)^2 = S^2$
Parabola	x_r, y_r, S_x, θ	$(y - y_r)^2 = 4S_x(x - x_r)^*$
Ellipse	$x_r, y_r, S_x, S_y, \theta$	$\frac{(y - y_r)^2}{S_y^2} + \frac{(x - x_r)^2}{S_x^2} = 1^*$

* Plus rotation by θ .

relatively few edges and large variations in parameters, the latter will be more efficient.

4. GENERALIZING THE HOUGH TRANSFORM

To generalize the Hough algorithm to non-analytic curves we define the following parameters for a generalized shape:

$$\mathbf{a} = \{y, \mathbf{s}, \theta\},$$

where $y = (x_r, y_r)$ is a reference origin for the shape, θ is its orientation, and $\mathbf{s} = (s_x, s_y)$ describes two orthogonal scale factors. As before, we will provide an algorithm for computing the best set of parameters \mathbf{a} for a given shape from edge pixel data. These parameters no longer have equal status. The reference origin location, y , is described in terms of a table of possible edge pixel orientations. The computation of the additional parameters \mathbf{s} and θ is then accomplished by straightforward transformations to this table. [To simplify the development slightly, and because of its practical significance, we will work with the four-dimensional sunspace $\mathbf{a} = (y, \mathbf{s}, \theta)$, where s is a scalar.]

In a sense this choice of parameters includes the previous analytic forms to which the Hough transform has been applied. Table 1 shows these relationships.

4.1 Earlier work: arbitrary shapes in binary edge images

Merlin and Farber⁽³⁾ showed how to use a Hough algorithm when the desired curves could not be described analytically. Each shape must have a specific reference point. Then we can use the following algorithm for a shape with boundary points B denoted by $\{\mathbf{x}_B\}$ which are relative to some reference origin y .

Merlin-Farber Hough algorithm: non-analytic curves with no gradient direction information $\mathbf{a} = y$. Form a two-dimensional accumulator array $A(\mathbf{a})$ initialized to zero. For each edge pixel x and each boundary point \mathbf{x}_B , compute \mathbf{a} such that $\mathbf{a} = x - \mathbf{x}_B$ and increment $A(\mathbf{a})$. Local maxima in $A(\mathbf{a})$ correspond to instances of the shape in the image.

Note that this is merely an efficient implementation of the convolution of the shape template where edge pixels are unity and others are zero with the corresponding image, i.e.,

$$A(\mathbf{x}) = T(\mathbf{x}) * S(\mathbf{x}) \quad (15)$$

where E is the binary edge image defined by

$$E(\mathbf{x}) = \begin{cases} 1 & \text{if } \mathbf{x} \text{ is an edge pixel} \\ 0 & \text{otherwise} \end{cases}$$

and $T(\mathbf{x})$ is the shape template consisting of ones where x is a boundary point and zeros otherwise, i.e.,

$$T(\mathbf{x}) = \begin{cases} 1 & \text{if } \mathbf{x} \text{ is in } B \\ 0 & \text{otherwise} \end{cases}$$

This result is due to Sklansky.⁽²⁰⁾

The Merlin-Farber algorithm is impractical for real image data. In an image with a multitude of edge pixels, there will be many false instances of the desired shape due to coincidental pixel arrangements. Nevertheless, it is the logical precursor to our generalized algorithm.

4.2 The generalization to arbitrary shapes

The key to generalizing the Hough algorithm to arbitrary shapes is the use of directional information. Directional information, besides making the algorithm faster, also greatly improves its accuracy. For example, if the directional information is not used in the circle detector, any significant group of edge points with quite different directions which lie on a circle will be detected. This can be appreciated by comparing Figs 2(a) and 2(b).

Consider for a moment the circular boundary detector with a fixed radius r_0 . Now for each gradient point x with direction ϕ , we need only increment a single point $\mathbf{x} + \mathbf{r}$. For the circle:

$$|\mathbf{r}| = r_0 \quad (16)$$

$$\text{Angle}(\mathbf{r}) = \phi(x) \quad (17)$$

Now suppose we have an arbitrary shape like the one shown in Fig. 6. Extending the idea of the circle detector with fixed radius to this case, for each point x on the boundary with gradient direction ϕ , we increment a point $\mathbf{a} = \mathbf{x} + \mathbf{r}$. The difference is that now $\mathbf{r} = \mathbf{a} - \mathbf{x}$ which, in general, will vary in magnitude and direction with different boundary points.




Fig. 6. Geometry for generalized Hough transform.

The fact that r varies in an arbitrary way means that the generalized Hough transform for an arbitrary shape is best represented by a table which we call the R -table.

4.3 The R -table

From the above discussion, we can see that the R -table is easily constructed by examining the boundary points of the shape. The construction of the table is accomplished as follows.

Algorithm for constructing an R -table. Choose a reference point y for the shape. For each boundary point x , compute $\phi(x)$ the gradient direction and $r = y - x$. Store r as a function of ϕ .

Notice that the mapping the table represents is vector-valued and, in general, an index ϕ may have many values of r . Table 2 shows the form of the R -table diagrammatically.

The R -table is used to detect instances of the shape S in an image in the following manner.

Generalized Hough algorithm for single shapes. For each edge pixel x in the image, increment all the corresponding points $x + r$ in the accumulator array A where r is a table entry indexed by ϕ , i.e., $r(\phi)$. Maxima in A correspond to possible instances of the shape S .

4.4 Examples

Some simple shapes are rotation-invariant, that is, the entries in the incrementation table are invariant functions of the gradient direction ϕ . Figure 7(a) shows an example for washers (or bagels). Here there are exactly two entries for each ϕ , one r units in the gradient direction and one R units in the direction opposite to the gradient direction. In another case the entries may be a simple function of ϕ . Figure 7(b)

Table 2. R -table format

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

shows such an example; hexagons. Irrespective of the orientation of the edge, the reference point locus is on a line of length l parallel to the edge pixel and $(3/2)l$ units away from it.

Another example is shown in Fig. 8. Here the points on the boundary of the shape are shown in Fig. 8(a). A reference point is selected and used to construct the R -table. Figure 8(b) shows a synthetic image of four different shapes and Fig. 8(c) shows the portion of the accumulator array for this image which has the correct values of orientation and scale. It is readily seen that edge points on the correct shape have incremented the same point in the accumulator array, whereas edge points on the other shapes have incremented disparate points.

4.5 R -table properties and the general notion of a shape

Up to this point we have considered shapes of fixed orientation and scale. Thus the accumulator array was two-dimensional in the reference point co-ordinates. To search for shapes of arbitrary orientation θ and scale s we add these two parameters to the shape description. The accumulator array now consists of four dimensions corresponding to the parameters (y, s, θ) . The R -table can also be used to increment this larger dimensional space since different orientations and scales correspond to easily-computed transformations of the table. Additionally, simple transformations to the R -table can also account for figure-ground reversals and changes of reference point.

We denote a particular R -table for a shape S by $R(\phi)$. R can be viewed as a multiply-vector-valued function. It is easy to see that simple transformations to this table will allow it to detect scaled or rotated instances of the same shape. For example if the shape is scaled by s and this transformation is denoted by T_s , then

$$T_s[R(\phi)] = sR(\phi) \quad (18)$$

i.e., all the vectors are scaled by s . Also, if the object is rotated by θ and this transformation is denoted by T_θ , then

$$T_\theta[R(\phi)] = \text{Rot}\{R[(\phi - \theta) \bmod 2\pi], \theta\} \quad (19)$$

i.e., all the indices are incremented by $-\theta$ modulo 2π , the appropriate vectors r are found, and then they are rotated by θ .

To appreciate that this is true, refer to Fig. 9. In this figure an edge pixel with orientation ϕ may be considered as corresponding to the boundary point x_A , in which case the reference point is y_A . Alternatively, the edge pixel may be considered as x_B on a rotated instance of the shape, in which case the reference point is at y_B which can be specified by translating r_A to x_B and rotating it through $+\Delta\theta$.

Figure-ground intensity reversals can also be taken into account via a simple R -table modification. The indices in the table are changed from ϕ to $(\phi + \pi) \bmod 2\pi$. Of course

$$T_{fg}\{T_{fg}[R(\phi)]\} = R(\phi)$$




Fig. 7. Simple examples using R -tables; (a) washers; (b) hexagons.




Fig. 8. An example. (a) Points on a shape used to encode R -table. (b) Image containing shape. (c) A plane through the accumulator array $A(x_r, y_r, S_0, \theta_0)$, where S_0 and θ_0 are appropriate for the shape in the image ($S_0 = 64, \theta_0 = 0$).




Fig. 9. Construction for visualizing the R -table transformation for a rotation by $\Delta\theta$. Point A can be viewed as: (1) on the shape (—), or (2) as point B on the shape (---), rotated by $\Delta\theta$. If (2) is used then the appropriate R is obtained by translating R_B to A and rotating it by $\Delta\theta$ as shown.

where T_{fg} denotes the figure-ground transformations.

Another property which will be useful in describing the composition of generalized Hough transforms is the change of reference point. If we want to choose a new reference point y' such that $y - y' = r$ then the modification to the R -table is given by $R(\phi) + r$, i.e. r is added to each vector in the table.

4.6 Using pairs of edges

We can also entertain the idea of using pairs of edge pixels to reduce the effort in parameter space. Using the R -table and the properties of the previous section, each edge pixel defines a surface in the four-dimensional accumulator space of $\mathbf{a} = (y, s, \theta)$. Two edge pixels at different orientations describe the same surface rotated by the same amount with respect to θ . Points where these two surfaces intersect (if any) correspond to possible parameters \mathbf{a} for the shape. Thus in a similar manner to Section 3.1, it is theoretically possible to use the two points in image space to reduce the locus in parameter space to a single point. However, the difficulties of finding the intersection points of the two surfaces in parameter space will make this approach unfeasible for most cases.

4.7 The Hough transform for composite shapes

Now suppose we have a composite shape S which has two subparts S_1 and S_2 . This shape can be detected by using the R -tables for S_1 and S_2 in a remarkably simple fashion. If y, y_1, y_2 are the reference points for shapes S, S_1 and S_2 respectively, we can compute $\mathbf{r}_1 = y - y_1$ and $\mathbf{r}_2 = y - y_2$. Then the composite generalized Hough transform $R_S(\phi)$ is given by

$$R_S(\phi) = [R_{S_1}(\phi) + \mathbf{r}_1] \cup [R_{S_2}(\phi) + \mathbf{r}_2] \quad (20)$$

which means that for each index value ϕ , \mathbf{r}_1 is added to $R_{S_1}(\phi)$, \mathbf{r}_2 is added to $R_{S_2}(\phi)$, and the union of these sets

is stored in $R_S(\phi)$. Equation 20 is very important as it represents a way of composing transforms.

In a similar manner we can define shapes as the difference between tables with common entries, i.e.,

$$R_S = R_{S_1} - R_{S_2} \quad (21)$$

means the shape S defined by S_1 with the common entries with S_2 deleted. The intersection operation is defined similarly. The primary use of the union operation is to detect shapes which are composites of simpler shapes. However, the difference operation also serves a useful function. Using it, R -tables which explicitly differentiate between two similar kinds of shapes can be constructed. An example would be differentiating between the washers and hexagons discussed earlier.

4.8 Building convolution templates

While equation (20) is one way of composing Hough transforms, it may not be the best way. This is because the choice of reference point can significantly affect the accuracy of the transform. Shapiro^(15,16,17) has shown this, emphasizing analytic forms. This is also graphically shown in Fig. 10. As the reference point becomes distant from the shape, small angular errors in ϕ can produce large errors in the vectors $R(\phi)$.

One solution to this problem is to use the table for each subshape with its own best reference point and to smooth the resultant accumulator array with a composite smoothing template. Recall that for the case of a single shape and isotropic errors (Section 2.2), convolving the accumulator array in this fashion was equivalent to taking account of the errors during the incrementation.

Where $h_i(y_i)$ denotes the smoothing template for reference point y_i of shape S_i the composite convolution template is given by




Fig. 10. Effects of changing reference point on errors.

$$H(y) = \sum_{i=1}^N h_i(y - y_i) \quad (22)$$

So finally, we have the following algorithm for the detection of a shape S which is the composite of subparts $S_1 \dots S_N$.

Generalized Hough algorithm for composite shapes. 1.

For each edge point with direction ϕ and for each value of scale s and orientation θ , increment the corresponding points $\mathbf{x} + \mathbf{r}$ in A where \mathbf{r} is in

$$R_s(\phi) = T_s \left\{ T_\theta \left[\bigcup_{k=1}^N R_{S_k}(\phi) \right] \right\}$$

2. Maxima in $A_s = A * H$ correspond to possible instances of the shape S . Figure 11 shows a simple example of how templates are combined.

If there are n edge pixels and M points in the error point spread function template, then the number of additions in the incrementation procedure is M . Thus this method might at first seem superior to the convolution method, which requires approximately $n^2 M$ additions and multiplications where $M < n^2$, the total number of pixels. However, the following heuristic is available for the convolution since A is typically very sparse. Compute

$$A_s(\mathbf{a}) \text{ only if } A(\mathbf{a}) > 0. \quad (23)$$

This in practice is very effective, although it may introduce errors if the appropriate index has a zero value and is surrounded by high values.

5. INCREMENTATION STRATEGIES

If we use the strategy of incrementing the accumulator array by unity, then the contents of the accumulator array are approximately proportional to the perimeter of the shape that is detectable in the image.

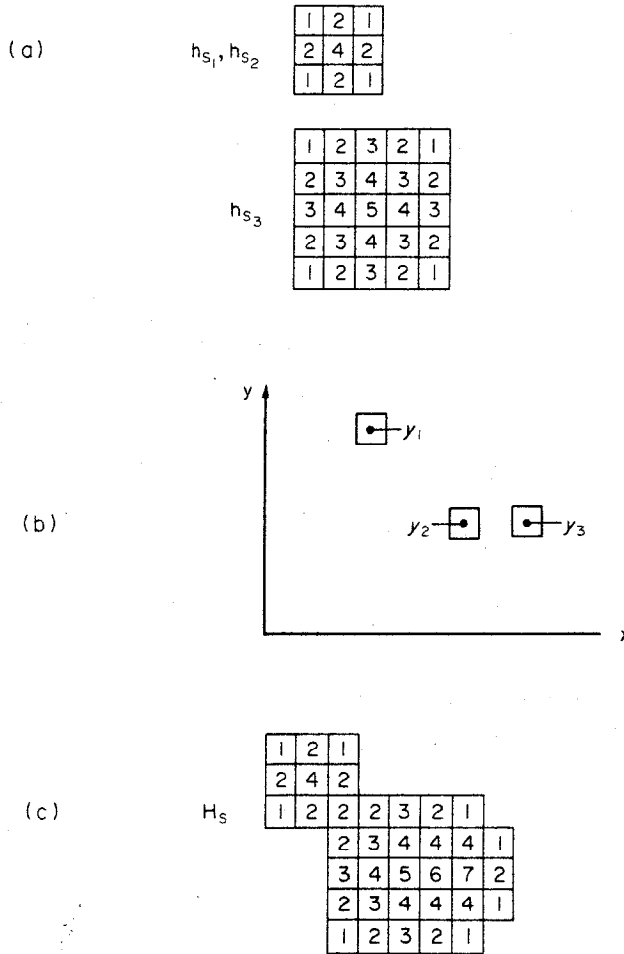


Fig. 11. Example of composite smoothing template construction. (a) Convolution templates for shapes S_1, S_2, S_3 . (b) Relationships between reference points $y_1, y_2,$ and y_3 in composite shape S . (c) Combined smoothing template H as a function of $h_1, h_2,$ and h_3 and $y_1, y_2,$ and y_3 .

This strategy is biased towards finding shapes where a large portion of the perimeter is detectable. Several different incrementation strategies are available, depending on the different quality of image data. If shorter, very prominent parts of the perimeter are detected, as might be the case in partially occluded objects, then an alternative strategy of incrementing by the gradient modulus value might be more successful, i.e.,

$$A(\mathbf{a}) := A(\mathbf{a}) + g(\mathbf{x}). \quad (24)$$

Of course the two strategies can be combined, e.g.,

$$A(\mathbf{a}) := A(\mathbf{a}) + g(\mathbf{x}) + c \quad (25)$$

where c is a constant.

Another possibility is the use of local curvature information in the incrementation function. Using this strategy, neighboring edge pixels are examined to calculate approximate curvature, K . This requires a more complicated operator than the edge operators we have considered, and complicates the table. Now along with each value of r the corresponding values of curvature must be stored. Then the incrementation

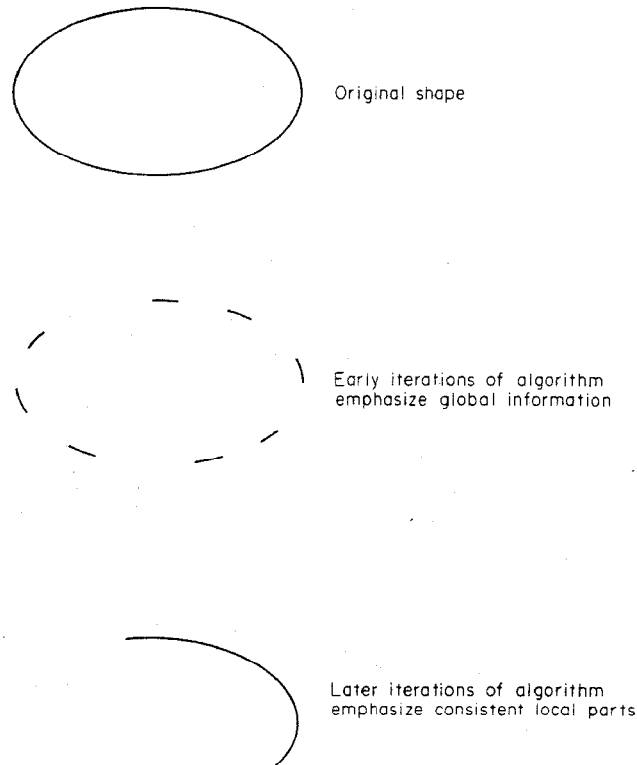


Fig. 12. Dynamic Hough transform.

weights 'informative' high local curvature edge pixels as follows:

$$A(\mathbf{a}) := A(\mathbf{a}) + K. \quad (26)$$

5.1 Weighting locally consistent information

Under certain circumstances we may want to weight local information that is consistent. For example, in searching for the boundary of that object, a connected set of edges conforming to the object may be more important than a set of unconnected edges. Figure 12 shows this example. Figure 12(a) might arise in situations with very noisy data. Figure 12(b) is an example where an object is occluded by another object. Wechsler and Sklansky,⁽⁶⁾ in the analytic formulation, successfully used the related strategy of increasing the incrementation factor if there were also neighboring edge pixels with the same edge direction. However, we would like to measure local consistency in parameter space.

A simple strategy for handling this case is to explicitly record the reference points for each edge pixel during a first pass. Then on a second pass edge pixels can increment by more than unity if neighboring edge pixels are incrementing the same reference point.

A more complicated strategy is to search for connected curve segments in image space which have compatible parameters. Such an algorithm, based on dynamic programming, is described in Ballard and Sklansky.⁽¹⁴⁾ The appropriate objective function for a

curve segment would be

$$h(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n) = \sum_{k=1}^n g(\mathbf{x}_k) + \sum_{k=1}^{n-1} q(\mathbf{x}_k, \mathbf{x}_{k+1}) \quad (27)$$

where

$$g(\mathbf{x}_k) = \text{the gradient magnitude} \quad (28)$$

and

$$q(\mathbf{x}_k, \mathbf{x}_{k+1}) = 0 \text{ if } |\phi(\mathbf{x}_k) - \phi(\mathbf{x}_{k+1})|_{\text{mod } \pi} \text{ is small and } -\infty \text{ otherwise} \quad (29)$$

In the dynamic programming algorithm, at each iteration step we can build longer compatible curves from all the edge points. Thus the incrementation function for a point \mathbf{x} would represent the longest compatible curve from that point. (If a longer curve cannot be built at any iteration, we can easily find this out.)

In a parallel implementation of this algorithm the contents of the accumulator array could be made to vary dynamically. Initially the contents would reflect global information, but with successive iterations the contents would be weighted in favor of consistent, local information.

5.2 More complex strategies

When searching for a composite object, different parts may have different importance. This is readily accommodated by associating a weight w_i with each table R_{S_i} so that each entry in R_{S_i} increments by a

factor w_i instead of unity.

The composite object may be searched for in a sequential manner. Applying the table sequentially could greatly improve the efficiency of the computations by limiting areas for subsequent suitable incrementations. Furthermore, standard methods^(21,22) could be used to stop the process once the shape had been located to the desired confidence level.

Even more complex strategies are possible wherein the process is integrated into a larger system. Here contextual information can be used to relegate all the previous operations including (a) building composite templates, (b) choosing weights, (c) choosing application sequences, and (d) adjusting weights in new contexts.

6. CONCLUSIONS

We have described a method for detecting instances of a shape S in an image which is a generalization of the Hough transform. This transform is a mapping from edge space to accumulator space such that instances of S produce local maxima in accumulator space. This mapping is conveniently described as a table of edge-orientation reference-point correspondence termed an R -table. This method has the following properties.

1. Scale changes, rotations, figure-ground reversals, and reference point translation of S can be accounted for by straightforward modifications to the R -table.

2. Given the boundary of the shape, its R -table can be easily constructed and requires a number of operations proportional to the number of boundary points.

3. Shapes are stored as canonical forms; instances of shapes are detected by knowing the transformation from the canonical form to the instance. If this transformation is not known then all plausible transformations must be tried.

4. If a shape S is viewed as a composite of several subparts $S_1 \dots S_n$ then the generalized Hough transform R -table for S can be simply constructed by combining the R -tables for $S_1 \dots S_n$.

5. A composite shape S may be efficiently detected in a sequential manner by adding the R -tables for the subparts S_i incrementally to the detection algorithm until a desired confidence level is reached.

6. The accumulator table values can be weighted in terms of locally consistent information.

7. The importance of a subshape S_i may be regulated by associating a weight w_i with the R -table.

8. Last but not least, the generalized Hough transform is a parallel algorithm.

Future work will be directed towards characterizing the computational efficiency of the algorithm and exploring its feasibility as a model of biological perception.

Acknowledgements - Portions of this paper benefitted substantially from discussions with Ken Sloan and Jerry Feldman. Special thanks go to R. Peet and P. Meeker for typing this manuscript. The work herein was supported by National Institutes of Health grant R23-HL21253-02.

REFERENCES

1. R. O. Duda and P. E. Hart, Use of the Hough transform to detect lines and curves in pictures, *Commun. Ass. Comput. Mach.* **15**, 11-15 (1975).
2. P. V. C. Hough, Method and means for recognizing complex patterns, U.S. Patent 3069654 (1962).
3. P. M. Merlin and D. J. Farber, A parallel mechanism for detecting curves in pictures, *IEEE Trans. Comput.* **C24**, 96-98 (1975).
4. F. O'Gorman and M. B. Clowes, Finding picture edges through collinearity of feature points, Proc. 3rd Int. Joint Conf. Artificial Intelligence, pp. 543-555 (1973).
5. C. Kimme, D. H. Ballard and J. Sklansky, Finding circles by an array of accumulators, *Commun. Ass. Comput. Mach.* **18**, 120-122 (1975).
6. H. Wechsler and J. Sklansky, Automatic detection of ribs in chest radiographs, *Pattern Recognition* **9**, 21-30 (1977).
7. S. A. Dudani and A. L. Luk, Locating straight-line edge segments on outdoor scenes, Proc. IEEE Computer Society on Pattern Recognition and Image Processing, Rensselaer Polytechnic Institute (1977).
8. C. L. Fennema and W. B. Thompson, Velocity determination in scenes containing several moving objects, Technical Report, Central Research Laboratory, Minnesota Mining and Manufacturing Co. St. Paul (1977).
9. J. R. Kender, Shape from texture: a brief overview and a new aggregation transform, Proc., DARPA Image Understanding Workshop, pp. 79-84. Pittsburgh, November (1978).
10. F. Attneave, Some informational aspects of visual perception, *Psychol. Rev.* **61**, 183-193 (1954).
11. D. Marr, Analyzing natural images: a computational theory of texture vision, MIT-AI-Technical Report 334, June (1975).
12. J. M. S. Prewitt, Object enhancement and extraction, *Picture Processing and Psychopictorics*, B. S. Lipkin and A. Rosenfeld, eds. Academic Press, New York (1970).
13. M. Hueckel, A local visual operator which recognizes edges and lines, *J. Ass. Comput. Mach.* **20**, 634-646 (1973).
14. D. H. Ballard and J. Sklansky, A ladder-structured decision tree for recognizing tumors in chest radiographs, *IEEE Trans. Comput.* **C25**, 503-513 (1976).
15. S. D. Shapiro, Properties of transforms for the detection of curves in noisy pictures, *Comput. Graphics Image Process.* **8**, 219-236 (1978).
16. S. D. Shapiro, Feature space transforms for curve detection, *Pattern Recognition* **10**, 129-143 (1978).
17. S. D. Shapiro, Generalization of the Hough transform for curve detection in noisy digital images, Proc. 4th Int. Joint Conf. Pattern Recognition, pp. 710-714. Kyoto, Japan, November (1978).
18. S. D. Shapiro, Transformation for the computer detection of curves in noisy pictures, *Comput. Graphics Image Process.* **4**, 328-338 (1975).
19. S. Tsuji and F. Matsumoto, Detection of elliptic and linear edges by searching two parameter spaces, Proc. 5th Int. Joint Conf. Artificial Intelligence, Vol. 2, pp. 700-705. Cambridge, MA, August (1977).
20. J. Sklansky, On the Hough technique for curve detection, *IEEE Trans. Comput.* **C27**, 923-926 (1978).
21. K. S. Fu, *Sequential Methods in Pattern Recognition and Machine Learning*. Academic Press, New York (1968).
22. R. Bolles, Verification vision with a programmable assembly system, Stanford AI Memo, AIM-275, December (1975).

APPENDIX. ANALYTIC HOUGH FOR PAIRS OF EDGE POINTS

To develop an explicit version of the Hough algorithm for ellipses using pairs of edge points, we consider the string-tied-at-two-ends parameterization of an ellipse:




Fig. A1. String-tied-at-both-ends parameterization of an ellipse.

$$(x-a)^2 + (y-b)^2 + (x-c)^2 + (y-d)^2 = l^2$$

where (a, b) and (c, d) are the ends and l is the length of the string, as shown in Fig. A1. Now if we have two edge points (x_1, y_1) and (x_2, y_2) with gradients ϕ_1 and ϕ_2 , the following equations result:

$$(x_1-a)^2 + (y_1-b)^2 + (x_1-c)^2 + (y_1-d)^2 = l^2 \quad (\text{A1})$$

$$(x_1-a) + (y_1-b)\phi_1 + (x_1-c) + (y_1-d)\phi_1 = 0 \quad (\text{A2})$$

$$(x_2-a)^2 + (y_2-b)^2 + (x_2-c)^2 + (y_2-d)^2 = l^2 \quad (\text{A3})$$

$$(x_2-a) + (y_2-b)\phi_2 + (x_2-c) + (y_2-d)\phi_2 = 0 \quad (\text{A4})$$

where in terms of the gradient direction

$$\phi = \tan\left(\phi - \frac{\pi}{2}\right) = \frac{dy}{dx}$$

From (A2):

$$a_1 = \phi_1[(y_1-b) + (y_1-d)] + 2x_1 - c$$

Substituting in (4):

$$x_2 - \phi_1[(y_1-b) + (y_1-d)] - 2x_1 + (y_2-b)\phi_2 + (y_2-d)\phi_2 + x_2 = 0.$$

Rearranging terms:

$$2x_2 - 2x_1 - 2\phi_1 y_2 + (\phi_1 - \phi_2)b + (\phi_1 - \phi_2)d = 0.$$

Now where:

$$S \equiv \phi_1 - \phi_2$$

and

$$K \equiv 2(x_2 - x_1 - \phi_1 y_1 + \phi_2 y_2)$$

and

$$t = \frac{K}{S}$$

we have

$$b = t - d. \quad (\text{A5})$$

Now we substitute for b in (2)

$$x_1 - a = -(x_1 - c) - \phi_1(2y_1 - t)$$

so that we have

$$c \equiv \eta - a \quad (\text{A6})$$

where $n = \phi_1(2y_1 - t)$

$$(x_1 - a)^2 + [y_1 - (t-d)]^2 + [x_1 - (\eta - a)]^2 + (y_1 - d)^2 = l^2 \quad (\text{A7})$$

$$(x_2 - a)^2 + [y_2 - (t-d)]^2 + [x_2 - (\eta - a)]^2 + (y_2 - d)^2 = l^2 \quad (\text{A8})$$

Thus our strategy for using two edge points is as follows:

Step 1: choose a .

Step 2: solve equations (5) and (6), a quadratic in d , for d .

Step 3: solve equation (2) for b and equation 3 for c .

Step 4: solve equation (1) for l .

Thus the vector $\mathbf{a} = (a, b, c, d, l)$ has been determined for a pair of edge pixels and can be used to increment the accumulator array.

About the Author—DANA H. BALLARD was born in Holyoke, MA, on October 15, 1946. He received the B.Sc. degree in aeronautics and astronautics from the Massachusetts Institute of Technology, Cambridge, in 1967. He received the M.S.E. degree in information and control engineering from the University of Michigan, Ann Arbor, in 1970 and the Ph.D. degree in information engineering from the University of California, Irvine, in 1974.

From 1968 to 1971 he was a Systems Analyst for Autonetics, Anaheim, Calif. Since 1971 he has been active in computer vision research. During the academic year 1974-1975, he was a Visiting Professor at the Laboratorio Biomediche Technologie, Rome, Italy. He is presently an Assistant Professor of Computer Science and Radiology at the University of Rochester, Rochester, N.Y. His current interests are in artificial intelligence and computer vision, particularly control strategies and geometric models, and their applications to biomedical image processing.