Ballard
780

TECH REPORT

780
10810

GENERALIZING THE HOUGH TRANSFORM
TO DETECT ARBITRARY SHAPES

Dana H. Ballard
Computer Science Department
University of Rochester

TR 55

October, 1979

The research described in this report was supported in part
by NIH Grant #R23-HL-21253-01, and in part by the Alfred P.
Sloan Foundation Grant # 78-4-15.

CARLSON LIBRARY

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 [Duda and Hart, 1972; Hough, 1962] and non-analytic curves [Merlin and Farber, 1975], 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 [O'Gorman and Clowes, 1973] and circles [Kimme et al., 1975] and parabolas [Wechsler and Sklansky, 1975]. The line detection case is the most well known and has been ingeniously exploited in several applications [Dudani and Lik, 1977; Fenema and Thompson, 1978; Kender, 1979].

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 a kind of universal transform which can be used to find arbitrarily complex shapes.

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 [Attneave, 1954] and that the image may be initially encoded as an "edge image," i.e., an image of local intensity or color gradients. Marr [1975] 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, 1970]) and the Hueckel operator [Hueckel, 1973], 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 coordinate 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.

Figure 1: Kinds of shapes detected with Generalized Hough.

- (a) simple shape
- (b) composite shape

The original algorithm by Hough [1962] did not use orientation information of the edge, and was considerably inferior to later work using the edge orientation for parametric curves [Kimme et al., 1975; Ballard and Sklansky, 1976; Wechsler and Sklansky, 1977]. Shapiro [1978] 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

To see how the Hough Transform works, let us suppose we are interested in detecting circular boundaries in an image. In rectangular coordinates, the equation for a circle is given by¹

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

Suppose also that the image has been partially 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 Figure 2a. This can be seen from Eq. 1 by treating x and y as fixed and letting, a , b , and r vary.

¹ Throughout the paper when a symbol is underlined it denotes a vector, but the same symbol without the underline denotes a scalar parameter. Thus, for example, we use $\underline{x}=(x,y)$ to denote a point.

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 would intersect at a common point.

Figure 2: (a) locus of parameters with no directional information;
(b) locus of parameters with directional information.

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 Figure 2b. 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.

The equation $\frac{df}{dx}(\underline{x}, \underline{a}) = 0$ introduces a term $\frac{dy}{dx}$ which is known since $\frac{dy}{dx} = \tan(\phi(\underline{x}) - \frac{\pi}{2})$ where $\phi(\underline{x})$ is the gradient direction. This suggests the following algorithm:

Hough Algorithm for Analytic Curves in Grey Level Images

For a specific curve $f(\underline{x}, \underline{a}) = 0$ with parameter vector

$\underline{a}$, form an array $A(\underline{a})$ initialized to zero. This array is termed an accumulator array. Then for each edge pixel $\underline{x}$, compute all $\underline{a}$ such that $f(\underline{x}, \underline{a}) = 0$ and $\frac{df}{d\underline{x}}(\underline{x}, \underline{a}) = 0$ and increment the corresponding accumulator array entries:

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

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

If only the equation $f(\underline{x}, \underline{a}) = 0$ 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 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 $A(\underline{a})$. Many sources of errors effect the computation of the parameter vector $\underline{a}$ so that in general many array locations in the vicinity of the ideal point $\underline{a}$ 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 [1975, 1978, 1978] 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 and the radius R . If we recognized these uncertainties as:

$$\begin{aligned} \phi(\underline{x}) \pm \Delta\phi \\ R \pm \Delta R(R) \end{aligned}$$

we would increment all values of $\underline{a}$ which fall within the shaded band of Figure 3. We let ΔR increase with R so that uncertainties are counted on a percentage basis.

Figure 3: Using convolution templates to compensate for errors.

Suppose we approximate this procedure by incrementing all values of $\underline{a}$ which fall inside the square domain centered about the nominal center shown in Figure 3, according to some point spread function, h .

After the first contributing pixel which increments center $\underline{a}_0$ has been taken into account, the new accumulator array contents A will be given by

$$A(\underline{a}) = h(\underline{a} - \underline{a}_0) \quad (\text{Eq. 2})$$

where $\underline{a} = (a_1, a_2, R)$ and $\underline{a}_0 = (a_{10}, a_{20}, 0)$.

If we include all the contributing pixels for that center, denoted by C , the accumulator is

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

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

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

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

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

Thus within the approximation of letting the square represent the shaded band shown in Figure 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 on page 8 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 an ubiquitous part of many everyday objects, appear as ellipses when viewed from a distant, oblique angle.

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

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 introduce an additional parameter for arbitrary orientations. then relax this requirement. For the moment, assume a and b are fixed. Then the equation of the ellipse is:

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

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

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

Differentiating with respect to X

$$\frac{2x}{a^2} + \frac{2y}{b^2} \frac{dy}{dx} = 0 \quad (\text{Eq. 8})$$

But $\frac{dy}{dx}$ is known from the edge pixel information! Let $\frac{dy}{dx} = \xi$, then from (4)

$$x^2 = \left(\frac{a^2}{b^2} \xi\right)^2 y^2 \quad (\text{Eq. 5})$$

So that substituting in (3)

$$\frac{y^2}{b^2} \left(1 + \frac{a^2}{b^2} \xi^2\right) = 1 \quad (\text{Eq. 10})$$

$$y = \pm \frac{b^2}{\sqrt{1 + \frac{a^2}{b^2} \xi^2}} \quad (\text{Eq. 11})$$

so that

$$x = \pm \frac{a^2}{\sqrt{1 + \frac{b^2}{a^2} \xi^2}} \quad (\text{Eq. 12})$$

so that finally given a, b, x, y, and $\frac{dy}{dx}$, we can determine x_0 and y_0 as:

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

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

The four solutions correspond to the four quadrants, as shown in Figure 5.

Figure 5: Four reference point solutions resolvable with gradient quadrant information.

The appropriate quadrant can be found by testing the signed differences dy and dx . So in ALGOL we would implement this as:

```
procedure HoughEllipse (integer  $x_{min}$   $x_{max}$ ,  $y_{min}$   $y_{max}$ ,  $\theta_{min}$   $\theta_{max}$ ,  
   $a_{min}$   $a_{max}$ ,  $b_{min}$   $b_{max}$ ,  $x$ ,  $y$ ,  $x_0$ ,  $y_0$ ,  $dx$ ,  $dy$ ; real angle,  $\xi$ ; integer  
  array A,  $\rho$  );
```

```
  begin;
```

```
  for  $x:=x_{min}$  step  $dx$  to  $x_{max}$  do
```

```
  for  $y:=y_{min}$  step  $dy$  to  $y_{max}$  do
```

```
    begin
```

```
       $dx:= p(x+1, y) - p(x, y)$ 
```

```
       $dy:= p(x, y+1) - p(x, y)$ 
```

```
      for  $a:=a_{min}$  step  $da$  until  $a_{max}$  do
```

```
      for  $b:=b_{min}$  step  $db$  until  $b_{max}$  do
```

```
      for  $\theta:=\theta_{min}$  step  $d\theta$  until  $\theta_{max}$  do
```

```
        begin
```

```
          angle:=  $\arctan(\frac{dy}{dx}) + \theta$ ;
```

```
           $\xi:= \tan(\text{angle})$ ;
```

$$dx := \text{SignX}(dX, dY) \sqrt{1 + \frac{a^2}{b^2} \xi^2}$$

$$dy := \text{SignY}(dX, dY) \sqrt{1 + \frac{b^2}{a^2} \xi^2}$$

Rotate-by-Theta(dx, dy);

$x_\theta := x + dx;$

$y_\theta := y + dy;$

$A(x_\theta, y_\theta, a, b) := A(x_\theta, y_\theta, \theta, a, b) + 1;$

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 . SignX and SignY are functions which return ± 1 depending on the quadrant determined by dX and dY .

3.1 Parameter Space Image Space Tradeoffs

Tsuji and Matsumoto [1977] 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 tradeoff 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 parameters. 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 technique will be more efficient. Whereas if there are 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:

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

where $\underline{y} = (x_r, y_r)$ is a reference origin for the shape, θ is its orientation, and $\underline{s} = (s_x, s_y)$ describes two orthogonal scale factors. As before, we will provide an algorithm for computing the best set of

parameters $\underline{a}$ for a given shape from edge pixel data. These parameters no longer have equal status. The reference origin location, $\underline{y}$, is described in terms of a table of possible edge pixel orientations. The computation of the additional parameters $\underline{s}$ 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 subspace $\underline{a} = (y, s, \theta)$, 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.

Table 1

4.1 Earlier Work: Arbitrary Shapes in Binary Edge Images

Merlin and Farber [1975] 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 $\{\underline{x}_B\}$ and reference point $\underline{y}$:

Merlin-Farber Hough Algorithm: Non-Analytic Curves
with no gradient direction information $\underline{a} = \{\underline{y}\}$

Form a two-dimensional accumulator array $A(\underline{a})$

initialized to zero. For each edge pixel $\underline{x}$ and boundary point $\underline{x}_B$, compute all $\underline{a}$ such that $\underline{a} - \underline{x} = \underline{y} - \underline{x}_B$ and increment $A(\underline{a})$. Local maxima in $A(\underline{a})$ 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.,

$$A(\underline{x}) = T(\underline{x}) * S(\underline{x}) \quad (\text{Eq. 15})$$

where E is the binary edge image defined by

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

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

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

This result is due to Sklansky [1977].

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 Figures 2a and 2b.

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

$$|\underline{r}| = r_0 \quad (\text{Eq. 16})$$

$$\text{Angle}(\underline{r}) = \phi(\underline{x}) \quad (\text{Eq. 17})$$

Figure 6: Geometry for Generalized Hough Transform.

Now suppose we have an arbitrary shape like the one shown in Figure 6. Extending the idea of the circle detector with fixed radius to this case, for each point $\underline{x}$ on the boundary with gradient direction , we

increment a point $\underline{a} = \underline{x} + \underline{r}$. The difference is that now $\underline{r} = \underline{a} - \underline{x}$ which, in general, will vary in magnitude and direction with different boundary points.

The fact that $\underline{r}$ 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 $\underline{y}$ for the shape. For each boundary point $\underline{x}$, compute $\phi(\underline{x})$ the gradient direction and $\underline{r} = \underline{y} - \underline{x}$. Quantize ϕ to one of $\lfloor \frac{2\pi}{n} \rfloor$ entries, and let $k \lfloor \frac{2\pi}{n} \rfloor$ be the quantized value where k is an integer in $[0, \dots, (n-1)]$. Store $\underline{r}$ as a function of k .

Henceforth we will use with the understanding that it is one of the n quantized values. Notice that the mapping the table represents is vector-valued and, in general, an index may have many values of $\underline{r}$. Figure 7 diagrams the form of the R-Table.

Figure 7: R-table format.

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 $\underline{x}$ in $\underline{P}$ such that $\phi(\underline{x}) = \phi_0$, increment all the corresponding points $\underline{x} + \underline{r}$ in the accumulator array A where $\underline{r}$ is a table entry indexed by ϕ_0 , i.e., $R(\phi_0)$. Maxima in A correspond to possible instances of the shape S .

4.4 Simple Examples: Rotation Invariant Tables

Some simple shapes are rotation-invariant, that is, the entries in the incrementation table are invariant functions of the gradient direction ϕ . Figure 8a 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. Figure 8b shows another such 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.

Figure 8: Simple examples using R-tables.

- (a) washers
- (b) hexagons

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 coordinates. To search for shapes of arbitrary orientation 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

$$T_s(R(\phi)) = sR(\phi) \quad (\text{Eq. 18})$$

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

$$T_\theta(R(\phi)) = \text{Rot}(R((\phi+\theta)\bmod 2\pi), \theta) \quad (\text{Eq. 19})$$

i.e., all the indices are incremented by θ modulo 2π and the vectors $\underline{r}$ are rotated by θ .

To appreciate that this is true, refer to Figure 9. In that figure an edge pixel with orientation ϕ may be considered as corresponding to the boundary point $\underline{x}_A$ in which case the reference point is $\underline{y}_A$. Alternatively, the edge pixel may be considered as $\underline{x}_B$ on a rotated instance of the shape in which case the reference point is at $\underline{y}_B$ which can be specified by translating $\underline{r}_B$ to $\underline{x}_A$ and rotating it through θ .

Figure 9: Modifications to the R-table to detect rotated shapes.

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

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

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 $\underline{y}'$ such that $\underline{y}-\underline{y}' = \underline{r}$ then the modification to the R-table is given by $R(\phi) + \underline{r}$, 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 $\underline{a} = (\underline{y}, s, \theta)$. 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 $\underline{a}$ 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 infeasible 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 $\underline{y}$, $\underline{y}_1$, $\underline{y}_2$ are the reference points for shapes S , S_1 and S_2 respectively, we can compute $\underline{r}_1 = \underline{y} - \underline{y}_1$ and $\underline{r}_2 = \underline{y} - \underline{y}_2$. Then the composite Generalized Hough Transform $R_S(\phi)$ is given by

$$R_S(\phi) = (R_{S_1}(\phi) + \underline{r}_1) \cdot (R_{S_2}(\phi) + \underline{r}_2) \quad (\text{Eq. 20})$$

which means that for each index value ϕ , r_1 is added to $R_{S_1}(\phi)$, 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.,

$$R_S = R_{S_1} \dot{-} R_{S_2} \quad (\text{Eq. 21})$$

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 hexagonals discussed earlier.

4.8 Building Convolution Templates

While Equation 28 is one way of composing Hough transforms, it is not the best way. This is because the choice of reference point can significantly effect the accuracy of the transform. Shapiro [1978] has shown this emphasizing analytic forms. This is also graphically shown in Figure 10. As the reference point becomes distant from the shape, small angular errors in ϕ can produce large errors in the

vectors $R(\phi)$.

Figure 10: Effects of changing reference point on errors.

The 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

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

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_0 and for each value of scale s and orientation θ , increment the corresponding points $\underline{x} + \underline{r}$ in A where $\underline{r}$ is in $R_S(\phi_0) = T_S(T_\theta(\bigcup_{k=1}^N R_{S_k}(\phi_0)))$
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.

Figure 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 .

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 seem superior at first to the convolution method which requires approximately n^2M 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

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

This in practice is very effective, although it may introduce errors if the appropriate index has value zero 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. 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 alternate strategy of incrementing by the gradient modulus value might be more successful, i.e.,

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

And of course the two strategies can be combined, e.g.,

$$A(\underline{a}) := A(\underline{a}) + g(\underline{x}) + c \quad (\text{Eq. 25})$$

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

$$A(\underline{a}) := A(\underline{a}) + K$$

(Eq. 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 12a might arise in situations with very noisy data. Figure 12b is an example where an object is occluded by another other object. Wechsler and Sklansky [1977], 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.

Figure 12: Different examples of a shape.

- a. global information
- b. consistent local information

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, 1976]. The appropriate objective function for a curve segment would be

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

where

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

and

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

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 $\underline{x}$ 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 drastically improve the efficiency of the computations by limiting areas for subsequent subtable incrementations. Furthermore, standard methods [Fu, 1968; Bolles, 1975] 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.

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:

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

where (a,b) and (c,d) are the ends and l is the length of the string, as shown in Figure A1.

Figure A1: String-tied-at-both-ends parametrization of an ellipse.

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:

$$(x_1-a)^2 + (y_1-b)^2 + (x_1-c)^2 + (y_1-d)^2 = l^2 \quad (\text{Eq. 1})$$

$$(x_1-a) + (y_1-b)\varphi_1 + (x_1-c) + (y_1-d)\varphi_1 = 0 \quad (\text{Eq. 2})$$

$$(x_2-a)^2 + (y_2-b)^2 + (x_2-c)^2 + (y_2-d)^2 = l^2 \quad (\text{Eq. 3})$$

$$(x_2-a) + (y_2-b)\varphi_2 + (x_2-c) + (y_2-d)\varphi_2 = 0 \quad (\text{Eq. 4})$$

where in terms of the gradient direction

$$\varphi = \tan(\phi - \pi/2) = \frac{dy}{dx}.$$

From (2): $a_1 = \varphi_1((y_1-b)+(y_1-d))+2x_1-c$

Substituting in (4):

$$x_2 - \varphi_1((y_1-b)+(y_1-d)) - 2x_1 + (y_2-b)\varphi_2 + (y_2-d)\varphi_2 + x_2 = 0$$

Rearranging terms: $2x_2 - 2x_1 - 2\varphi_1 y_1 + (\varphi_1 - \varphi_2)b + (\varphi_1 - \varphi_2)d + 2\varphi_2 y_2 = 0$

Now where: $S \equiv \varphi_1 - \varphi_2$

and $K \equiv 2(x_2 - x_1 - \varphi_1 y_1 + \varphi_2 y_2)$

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

we have

$$b = t - d \quad (\text{Eq. 5})$$

Now we substitute for b in (2)

$$x_1 - a = -(x_1 - c) - \varphi_1(2y_1 - t)$$

so that we have

$$c = n - a \quad (\text{Eq. 6})$$

where $n \equiv \varphi_1(2y_1 - t)$

$$(x_1 - a)^2 + (y_1 - (t - d))^2 + (x_1 - (n - a))^2 + (y_1 - d)^2 = \ell^2 \quad (\text{Eq. 7})$$

$$(x_2 - a)^2 + (y_2 - (t - d))^2 + (x_2 - (n - a))^2 + (y_2 - d)^2 = \ell^2 \quad (\text{Eq. 8})$$

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

Step 1: choose a

Step 2: solve Eq. 5-6, a quadratic in d , for d

Step 3: solve Eq. 2 for b and Eq. 3 for c

Step 4: solve Eq. 1 for ℓ .

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

Table 1

<i>Analytic Form</i>	<i>Parameters</i>	<i>Equation</i>
<i>Line</i>	s, θ	$x \cos \theta + y \sin \theta = s$
<i>Circle</i>	x_r, y_r, s	$(x - x_r)^2 + (y - y_r)^2 = s^2$
<i>Parabola</i>	x_r, y_r, s_x, θ	$(y - y_r)^2 = 4 s_x (x - x_r)^*$
<i>Ellipse</i>	$x_r, y_r, s_x, s_y, \theta$	$\frac{(y - y_r)^2}{s_y^2} + \frac{(x - x_r)^2}{s_x^2} = 1^*$

* plus rotation by θ

Analytic Curves

described in terms of the generalized shape parameters $x_r, y_r, s_x, s_y, \theta$

Figure 1: Kinds of shapes detected with Generalized Hough.

(a) simple shape

(b) composite shape

a.

b.

Figure 2: (a) locus of parameters with no directional information;
(b) locus of parameters with directional information.

Figure 3: Using convolution templates to compensate for errors.

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

Figure 5: Four reference point solutions resolvable with gradient quadrant information.

Figure 6: Geometry for Generalized Hough Transform.

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

Figure 7: R-table format.

Figure 8: Simple examples using R-tables.

(a) washers

(b) hexagons

Figure 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) then the appropriate R is obtained by translating R_B to A and rotating it by $\Delta\theta$ as shown.

Figure 10: Effects of changing reference point on errors.

Original shape

Early iterations of algorithm
emphasize global information

Later iterations of algorithm
emphasize consistent local parts

Figure 12: Dynamic Hough Transform

a:

h_{s_1}, h_{s_2}

1	2	1
2	4	2
1	2	1

h_{s_3}

1	2	3	2	1
2	3	4	3	2
3	4	5	4	3
2	3	4	3	2
1	2	3	2	1

b:

c:

H_S

1	2	1					
2	4	2					
1	2	2	2	3	2	1	
		2	3	4	4	4	1
		3	4	5	6	7	2
		2	3	4	4	4	1
		1	2	3	2	1	

Figure 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 .

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

REFERENCES

- Attneave, F., "Some informational aspects of visual perception," Psychol. Rev. 61, 1954.
- Ballard, D. H. and Sklansky, J., "A ladder-structured decision tree for recognizing tumors in chest radiographs," IEEE Transactions on Computers, Vol C-25, 1976, pp. 503-513.
- Bolles, R., "Verification vision with a programmable assembly system," Stanford AI Memo, AIM-275, Dec. 1975.
- Duda, R.O. and Hart, P.E., "Use of the Hough transform to detect lines and curves in pictures," Comm.ACM, 15, 1, Jan. 1972, pp. 11-15.
- Dudani, S.A. and A.L. Luk, "Locating straight-line edge segments on outdoor scenes," Proceedings, IEEE Computer Society on Pattern Recognition and Image Processing, Rensselaer Polytechnic Institute, 1977.
- Fennema, C.L. and W.B. Thompson, "Velocity determination in scenes containing several moving objects," Technical Report, Central Research Laboratory, Minnesota Mining and Manufacturing Company, St. Paul, 1977.
- Fu, K.S., Sequential Methods in Pattern Recognition and Machine Learning, Academic Press, 1968.
- Hueckel, M., "A local visual operator which recognizes edges and lines," JACM, Vol. 20, No. 4, Oct. 1973.
- Hough, P.V.C., "Method and means for recognizing complex patterns," U.S. Patent 3,069,654, 1962.
- Kender, J.R., "Shape from texture: a brief overview and a new aggregation transform," Proceedings, DARPA Image Understanding Workshop, Pittsburgh, Pennsylvania, November, 1978, 79-84.
- Kimme, C., Ballard, D.H., and Sklansky, J., "Finding circles by an array of accumulators," Comm.ACM, 18, 1, Feb. 1975, pp.120-122.
- Marr, D., "Analyzing natural images: a computational theory of texture vision," MIT-AI-Tech. Repot 334, June, 1975.
- Merlin, P.M. and Farber, D.J., "A parallel mechanism for detecting curves in pictures," IEEE TRANS. on Computers, C24, Jan. 1975, pp. 96-98.
- O'Gorman, F., and Clowes, M.B., "Finding picture edges through collinearity of feature points," Proc. Third IJCAI, 1973, pp. 543-555.

- Prewitt, J.M.S., "Object Enhancement and Extraction," in Picture Processing and Psychopictorics, B.S. Lipkin and A. Rosenfeld, Eds., Academic Press, 1970.
- Sklansky, J., "On the Hough technique for curve detection," IEEE Trans. on Computers, July, 1977.
- Shapiro, S.D., "Transformation for the computer detection of curves in noisy pictures," Computer Graphics and Image Processing, 4, pp. 328-338, 1975.
- Shapiro, S.D., "Properties of transforms for the detection of curves in noisy pictures," Computer Graphics and Image Processing, 8, 2, pp. 219-236, October, 1978.
- Shapiro, S.D., "Feature space transforms for curve detection," Pattern Recognition 10, 3, pp. 129-143, 1978.
- Shapiro, S.D., "Generalization of the Hough transform for curve detection in noisy digital images," Proc.. 4th IJ CPR, Kyoto, Japan, Nov. 1978, pp. 710-714.
- Tsuji, S. and Matsumoto, F., "Detection of elliptic and linear edges by searching two parameter spaces," Proc. 5th IJCAI, Vol 2, Cambridge, Mass., Aug. 1977, pp. 700-705.
- Wechsler, H. and Sklansky, J., "Automatic detection of ribs in chest radiographs," Pattern Recognition, 9, Jan. 1977, pp. 21-30.