

Graphics and
Image Processing

W. Newman
Editor

Use of the Hough Transformation To Detect Lines and Curves in Pictures

Richard O. Duda and Peter E. Hart
Stanford Research Institute,
Menlo Park, California

Hough has proposed an interesting and computationally efficient procedure for detecting lines in pictures. This paper points out that the use of angle-radius rather than slope-intercept parameters simplifies the computation further. It also shows how

1. Introduction

A recurring problem in computer picture processing is the detection of straight lines in digitized images. In

the method can be used for more general curve fitting, and gives alternative interpretations that explain the source of its efficiency.

Key Words and Phrases: picture processing, pattern recognition, line detection, curve detection, colinear points, point-line transformation, Hough transformation

CR Categories: 3.63

Copyright (~ 1972, Association for Computing Machinery, Inc. General permission to republish, but not for profit, all or part of this material is granted, provided that reference is made to this publication, to its date of issue, and to the fact that reprinting privileges were granted by permission of the Association for Computing Machinery.

This work was supported by the Advanced Research Projects Agency and the National Aeronautics and Space Administration under Contract NAS12-2221.

the simplest case, the picture contains a number of discrete, black figure points lying on a white background.

The problem is to detect the presence of groups of colinear or almost colinear figure points. It is clear that the problem can be solved to any desired degree of accuracy by testing the lines formed by all pairs of points. However, the computation required for n points is approximately proportional to n^2 , and may be prohibitive for large n .

Rosenfeld [1] has described an ingenious method due to Hough [2] for replacing the original problem of finding colinear points by a mathematically equivalent problem of finding concurrent lines. This method involves transforming each of the figure points into a straight line in a *parameter space*. The parameter space is defined by the parametric representation used to describe lines in the picture plane. Hough chose to use the familiar slope-intercept parameters, and thus his parameter space was the two-dimensional slope-intercept plane. Unfortunately, both the slope and the intercept are unbounded, which complicates the application of the technique. In this note we suggest an alternative parameterization that eliminates this problem. We also give two alternative interpretations of Hough's method, one of which reveals plainly the source of its efficiency. Finally, we show how the method can be extended to find more general classes of curves in pictures.

2. Fundamentals

The set of all straight lines in the picture plane constitutes a two-parameter family. If we fix a parameteriza-

tion for the family, then an arbitrary straight line can be represented by a single point in the parameter space. For reasons that become obvious, we prefer the so-called *normalparameterization*. As illustrated in Figure 1, this parameterization specifies a straight line by the angle θ of its normal and its algebraic distance p from the origin. The equation of a line corresponding to this geometry is

$$x \cos \theta + y \sin \theta = p.$$

If we restrict θ to the interval $[0, \pi]$, then the normal parameters for a line are unique. With this restriction, every line in the x - y plane corresponds to a unique point in the θ - p plane.

Suppose now, that we have some set $\{ (x_a, y_a), \dots, (x_n, y_n) \}$ of n figure points and we want to find a set of straight lines that fit them. We transform the points (x_i, y_i) into the sinusoidal curves in the θ - p plane defined by

$$p = x_i \cos \theta + y_i \sin \theta. \quad (1)$$

It is easy to show that the curves corresponding to colinear figure points have a common point of intersection. This point in the θ - p plane, say (θ_0, p_0) , defines the line passing through the colinear points. Thus the problem of detecting colinear points can be converted to the problem of finding concurrent curves.

A dual property of the point-to-curve transformation can also be established. Suppose we have a set $\{ (\theta_1, p_1), \dots, (\theta_k, p_k) \}$ of points in the θ - p plane, all lying on the curve

$$p = X \cos \theta + y \sin \theta.$$

Then it is easy to show that all these points correspond to lines in the x - y plane passing through the point (x_0, y_0) . We can summarize these interesting properties of the point-to-curve transformation as follows:

Property 1. A point in the picture plane corresponds to a sinusoidal curve in the parameter plane.

Property 2. A point in the parameter plane corresponds to a straight line in the picture plane.

Fig. 1. The normal parameters for a line.

x

Fig. 2. Projection of colinear points onto a line.

Property 3. Points lying on the same straight line in the picture plane correspond to curves through a common point in the parameter plane.

Property 4. Points lying on the same curve in the parameter plane correspond to lines through the same point in the picture plane.

In Section 3 we apply these results to the problem of detecting colinear points in the picture plane and show how significant computational economies can be realized in certain situations.

3. Applications and Alternative Interpretations

Suppose we map all of the points in the picture plane into their corresponding curves in the parameter plane. In general, these n curves will intersect in $n(n-1)/2$ points corresponding to the lines between all pairs of

{x2, Y2 }

12

Communications
of
the ACM

January 1972
Volume 15
Number 1

Page 3

figure points. Exactly colinear subsets of figure points can be found, at least in principle, by finding coincident points of intersection in the parameter plane. Unfortunately, this approach is essentially exhaustive, and the computation required grows quadratically with the number of picture points.

When it is not necessary to determine the lines exactly, the computational burden can be reduced considerably. Following Hough's basic proposal, we specify the acceptable error in θ and p and quantize the θ - p plane into a quadrated grid. This quantization can be confined to the region $0 \leq \theta < 2\pi$, $-R \leq p \leq R$, where R is the size of the retina, since points outside this rectangle correspond to lines in the picture plane that do not cross the retina. The quantized region is

2: Points lying on the same curve in the θ - p plane correspond to lines through the same point in the picture plane. When the curve corresponding to figure point (x_i, y_i) is "added" to the accumulator, we are really computing and recording the parameters of the d : lines in the picture plane passing through (x_i, y_i) , and because θ is quantized, these are "all the lines in the plane" passing through (x_i, y_i) . Should a given parameter pair ever recur as a result of computing the d : lines through some other figure point, the recurrence will be reflected in an increased count in the appropriate accumulator cell. Roughly speaking, then, for each figure point the quantized transform method considers only the set of all d : lines through that point,

treated as a two-dimensional array of accumulators. For each point (x_i, y_i) in the picture plane, the corresponding curve given by (1) is entered in the array by incrementing the count in each cell along the curve. Thus, a given cell in the two-dimensional accumulator eventually records the total number of curves passing through it. After all figure points have been treated, the array is inspected to find cells having high counts. If the count in a given cell (O_i, p_i) is k , then precisely k figure points lie (to within quantization error) along the line whose normal parameters are (O_i, p_i) .

An alternative interpretation of the point-curve transformation can be obtained by recognizing that the p computed by (1) locates the projection of the point (x_i, y_i) onto a line through the origin with slope angle θ . Thus, if a number of figure points lie close to some line l , their projections onto the line normal to l are nearly coincident (see Figure 2). A given column in the $O-p$ accumulator array is just a histogram for these projections, so a high count in a given cell clearly corresponds to a nearly colinear subset of figure points. A variation of this approach was used by Griffith [3] to find long lines in a picture.

Let us investigate how the computation required by the accumulator implementation varies with the number of figure points. To be more specific about the quantization, suppose that we restrict our attention to d_θ values of θ uniformly spaced in the interval $[0, 7\pi/4)$. Suppose further that the p axis in the interval $[-R, R]$ is quantized into d_p cells. For each figure point (x_i, y_i) , we use (1) to compute the d_p different values of p corresponding to the d_θ possible values of the independent variable θ . Since there are n figure points, we need to carry out this computation nd_p times. When these computations are complete, the $d_\theta d_p$ cells of the two-dimensional accumulator are inspected to find high counts. Thus the computation required grows linearly with the number of figure points. Clearly, when n is large compared to $d_\theta d_p$, this approach is preferable to an exhaustive procedure that requires considering the lines between all $n(n-1)/2$ pairs of figure points.

whereas more exhaustive methods consider all $(n-1)$ lines between the given point and all other figure points.

4. Example

The following example illustrates some of the features of the transform approach. Figure 3 (a) shows a television monitor view of a box, and Figure 3 (b) shows a digitized version of that view. A simple differencing operation locates significant intensity changes and produces the binary picture shown in Figure 3(c). This 120 X 120 picture contains many nearly colinear figure points that can be fit well by a few straight lines.

Sampling θ at $d_\theta = 9$ twenty-degree increments in 0 and, quantizing p into $d_p = 86$ two-element cells, we obtain the two-dimensional accumulator array shown in Table I. If the array entry at (O_i, p_i) is k_i , then k_i figure points lie on parallel lines for which $\theta = O_i$, and p_i lies between p_{i-1} and p_i . When many points are nearly colinear, the entry for the line that fits them best is large. The largest entry in Table I occurs at $(0^\circ, -5)$ and corresponds to the middle vertical edge of the box. The nine circled entries correspond to locally maximum values that exceed the arbitrary threshold of 35. The corresponding nine groups of nearly colinear figure points are shown in Figure 3(d). In this example, it happens that every group corresponds to some physically meaningful line in the picture. However, two significant lines on the top of the box were not found--one, because it contained very few points, and the other, because it fell between the lines at $\theta = 80^\circ$ and $\theta = 100^\circ$. The 20° angular quantization interval was chosen to keep the accumulator array small. Clearly, we were fortunate to have found as many lines as we did, and a smaller quantization interval would have to be used in practice.

A few remarks concerning some limitations of the transform approach are in order. First, the results are sensitive to the quantization of both θ and p . Finer

Another alternative interpretation exposes the source of this efficiency. Consider again Property 4 in Section

13

quantization gives better resolution, but increases the computation time and exposes the problem of clustering entries corresponding to nearly colinear points. Second,

Communications
of
the ACM

January 1972
Volume 15
Number 1

Page 4

the technique finds colinear points without regard to contiguity. Thus the position of a best-fit line can be distorted by the presence of unrelated figure points in another part of the picture. A related problem is that of meaningless groups of colinear points being detected. In our example, a false line would be detected if the threshold were reduced from 35 to 24, the value needed to detect the top left-hand edge of the box.

The transform approach does successfully find groups of colinear or nearly colinear figure points. If the minimum size of a significant group is known, all such groups can be detected. If additional properties such as contiguity are known, they can be used to reject meaningless results. In general, the transform approach should be viewed primarily as a computationally efficient way of accomplishing a conceptually simple step in scene analysis.

5. Extensions and Conclusions

The transform method can be generalized and specialized in several ways. We note immediately that any parametrization of the family of straight lines can be used. As we have mentioned, Hough used the slope-intercept parameterization. However, this parameterization has the disadvantage of being sensitive to the choice of coordinate axes in the picture plane. If several figure points lie on a nearly vertical line, for example, both the slope and the intercept may be arbitrarily large. Thus

Fig. 3 An illustrative example.

1

1

(a) Monitor

(b) Digitized

the entire two-dimensional parameter plane must be considered. As Rosenfeld [1] has pointed out, one could do the entire problem twice, interchanging the x- and y-axes, but this would introduce additional complications. The normal parametrization avoids these disadvantages, fundamentally for the same reasons that make it useful in integral geometry: It allows us to place an invariant measure on the set of all straight lines.

An important special use of the transform method is to detect the occurrence of figure points lying on a straight line and possessing some specified property. For example, suppose we want to find whether a significant number of figure points lie on a line through the point (x_0, y_0) in the picture plane. As we have seen from Property 4, the normal coordinates of any such line must lie on (or, in practice at least near) the curve $n = x_0 \cos \theta + y_0 \sin \theta$. Hence, the transform process can be carried out in the usual way, but attention can be restricted to the region of the θ - p plane near this curve. If we find a cell with count k near this curve, then we are assured that k figure points lie on a line passing (nearly) through the point (x_0, y_0) . Similarly, suppose we are interested only in lines having a given direction, say θ_0 . Again, we carry out the process in the usual way, but restrict our attention to a subset of the θ - p plane in the vicinity of $\theta = \theta_0$.

It is clear that the general transform approach can be

(c) Gradient

ii,

(d) Lines

Table I. AccumulatorArray for Figure 3(c)

\backslash	$\overset{0}{\sim\sim\sim}0:$	20: 40: 60: 80 ° zoo: 120: 140:160:	\circ	$\overset{0}{\sim}0=$	20 = 40: 60:	80: 100 = 120 ° 140:160 °
--------------	-------------------------------	-------------------------------------	---------	-----------------------	--------------	---------------------------

formed into a surface in the a - b - c parameter space defined by

$$(x_i - a)^2 + (y_i - b)^2 = c^2.$$

in this example, then, each figure point will be transformed into a right circular cone in a three-dimensional parameter space. If the cones corresponding to many figure points intersect at a single point, say the point (a_0, b_0, c_0) , then all the figure points lie on the circle defined by those three parameters. As in the preceding case of straight lines, no saving is effected if the entire process is performed analytically. However, the process can be implemented efficiently by using a three-dimen-

parameter space and the computation grown exponentially with the number of parameters.

Received November 1970, revised April 1971

References

1. Rosenfeld, A. *Picture Processing by Computer*. Academic Press, New York, 1969.
2. Hough, P.V.C. Method and means for recognizing complex patterns. U.S. Patent 3,069,654, Dec. 18, 1962.
3. Griffith, A.K. Computer recognition of prismatic solids. Ph.D. Th., Dep. of Math., MIT, June 1970.