THE PERFORMANCE OF THE GENERALISED HOUGH TRANSFORM: CONCAVITIES, AMBIGUITIES AND POSITIONAL ACCURACY

E R Davies
Machine Vision Group, Department of Physics
Royal Holloway & Bedford New College
Egham Hill, Egham, Surrey TW20.0EX, UK

ABSTRACT

This paper studies the use of the generalised Hough transform to locate objects possessing concavities. When locating objects with internal symmetries, ambiguities can arise which may not be resolvable in the presence of occlusions. On the other hand concavities can help to make object location more accurate. The boundary orientation distribution can be used to analyse the situation: it shows that enhanced accuracy of location along one axis may result in reduced accuracy in a perpendicular direction. Finally, accuracy of location is independent of the position of the object localisation point.

I INTRODUCTION

Although the Hough transform was originally devised for the detection of straight lines as long ago as 1962 (Hough 1962), it only came into wide use in the image processing community after it was 're-discovered' by Rosenfeld in 1969 and further developed by Duda and Hart in 1972. Subsequent work applied the technique to the detection of circles, at the same time making it more efficient by showing how locally available edge orientation information could be used to cut down the number of votes accumulated in parameter space (Kimme et al 1975). Later, the technique was applied to other specific types of curve such as ellipses and parabolas before Ballard finally generalised it so that it could be applied to the detection of arbitrary shapes (Ballard 1981). The resulting 'generalised Hough transform' (GHT) retained the facility for making use of local edge orientation information, and is thus a highly efficient procedure.

More recently, Davies has found that one advantage of using the GHT rather than the basic Hough transform to detect straight edges is that this enables objects such as squares and rectangles to be detected directly - i.e. without further high-level processing to deduce the presence of the shape (Davies 1986). He has also shown how to ensure that optimal sensitivity is attained while using the technique to detect objects possessing straight edges (Davies 1987b). Finally, he has demonstrated that the GHT is exceptionally robust and is therefore particularly suitable for industrial (e.g. automated inspection) applications (Davies 1984).

The author is grateful to the UK Science and Engineering Research Council for financial support during the course of this research.

This paper starts with the observation that a problem arises when the GHT is used to detect objects possessing concavities. In principle this can lead to ambiguities in the location of such objects, and it was felt necessary to investigate the problem closely in order to find how its effects could be minimised.

In section II we briefly describe the operation of the GHT. Then in section III we consider the problem posed by concavities, following this in section IV by a discussion of the effects of symmetries. In sections V and VI we investigate how the accuracy with which an object can be located is affected by its shape and by errors in the estimation of edge orientation.

II THE GENERALISED HOUGH TRANSFORM

The action of the GHT may be summarised as follows. First, a Sobel or similar edge enhancement operator is applied to the image, and the resulting intensity gradient image is thresholded to find the locations of the most significant edge pixels. The location and orientation of each edge pixel are then used to estimate the position of a localisation point L within every object of a specific type supposed to be present in the original image. All such estimated locations of L are accumulated in a parameter space which is, for the GHT, congruent to image space. Finally, peaks are sought in parameter space which indicate the presence and position of objects of the chosen shape. For objects of other shapes, the points accumulated in parameter space are essentially incoherent and do not focus on peaks: such objects effectively contribute noise in parameter space, and are unlikely to interfere with the process of locating the chosen type of object (Ballard 1981).

The precise way in which the position of L is estimated after locating an edge pixel can in simple cases be analytic. For example, in the case of a circle, L is found merely by moving a distance equal to R along the edge normal. For complex curves, if the edge normal orientation is θ, we move a distance $R(\theta)$ in a direction $\varphi(\theta)$ from the edge location, the values of R and φ being obtained from a lookup table.

As outlined so far, the orientation of the curve has to be known in advance. However, if object orientation is able to vary, we merely adopt the strategy of augmenting parameter space in another dimension, each plane in parameter space then being used to detect the object in one of its possible orientations. In what follows we simplify the discussion by assuming the orientation of the object is known. Finally, because of possible size and shape variations, it will often be preferable to save computation by using the GHT to look for

AVC 1987 doi:10.5244/C.1.44

features rather than for whole objects: suitable features include holes, corners, lines, circular arcs and so on. Again, we ignore such complications in what follows.

III THE EFFECTS OF CONCAVITIES ON OBJECT DETECTION BY THE GHT

Before proceeding to study the effects of concavities on object detection, it is relevant to outline the problems that arise when objects possess straight edges. Basically, for a simple convex object, each edge pixel gives rise to a single vote in parameter space. As described in section II, this vote will be placed at the estimated position of the object localisation point L. However, when an object possesses a straight edge, this strategy severely limits the number of votes accumulated at L by the pixels on a particular straight edge - typically only one or two votes being accumulated: hence sensitivity is markedly reduced. Davies showed that this problem could only be eliminated by accumulating a line of points in parameter space for each straight edge pixel, the appropriate length of the line being equal to the length of the straight portion of the boundary in the object being detected, and the direction being that of the straight portion of the object boundary (Davies 1987b). At a particular edge pixel, this direction is estimated by taking the local edge orientation.

The rationale for this parametrisation of the straight line is as follows. For a given edge pixel whose direction indicates it is on the straight portion of the object boundary, we accumulate a vote at all possible positions of the localisation point which would confirm this as being a valid straight edge pixel. Though this rationale might seem clever but somewhat ad hoc, Davies was able to confirm its value on the basis that the GHT has to approximate a spatial matched filter if sensitivity is to be optimised (Davies 1987b).

The situation when concavities arise is closely related to that for straight edges. In particular, when a single concavity is present, each of the edge pixels within the concavity has the same orientation as one other edge pixel on the object boundary. (But note that certain concavities can act as double concavities in this respect - see Figure 1.) This means that, for a particular range of edge orientations, it is not known locally at the edge pixel whether it is on that portion of the boundary within the concavity or elsewhere. Thus there is an ambiguity in our local knowledge of the position of the edge pixel relative to the localisation point - just as there is a multiple ambiguity in the case of pixels on a straight edge - and there is no choice but to accumulate two votes in parameter space for such edge pixels. One could argue that the case of a concavity is more extreme than that of a straight edge, but in fact the straight edge gives rise to a multiple ambiguity which involves accumulating a great many points in parameter space for each edge pixel, and this is computationally more serious.

For more complex types of concavity - multiple holes, sawtooth edges and so on - the local ambiguity increases, and with it the number of votes that have to be accumulated per edge pixel, for at least a proportion of local edge orientations. Indeed, for spirals of many sorts and all cases of holes, at least two votes have to be accumulated per edge pixel for all possible edge orientations.

Though a potential problem has arisen with concavities, we see that the rationale that enabled us to cope with straight edges also permits us to cope with concavities, and at the same time to attain the maximum sensitivity. It is clear that sensitivity is maximised since every edge pixel contributes an equal amount to the peak, at L, in parameter space: thus all the available 'signal' is utilised.

IV SYMMETRIES THAT ARISE FROM CONCAVITIES

Though we apparently solved the concavity problem successfully in the previous section, this was achieved by accumulating additional points in parameter space. Such points lead to additional 'clutter' in parameter space which could conceivably focus into unwanted peaks. This means that it is necessary to examine whether the additional points that are accumulated could make phantom objects appear. Unfortunately, there are situations when this can happen. In particular, if any object possesses symmetries between separate parts of its shape, then spurious peaks can appear in parameter space. Consider the v-shaped object shown in Figure 2. There is a 2-fold translational symmetry

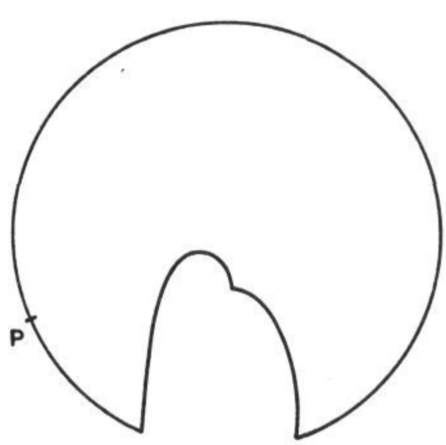

Figure 1 A shape with a structured concavity

This diagram shows a concavity which must be classed as a 'double' concavity: thus there are two points in the concavity with the same orientation as point P on the convex part of the boundary.

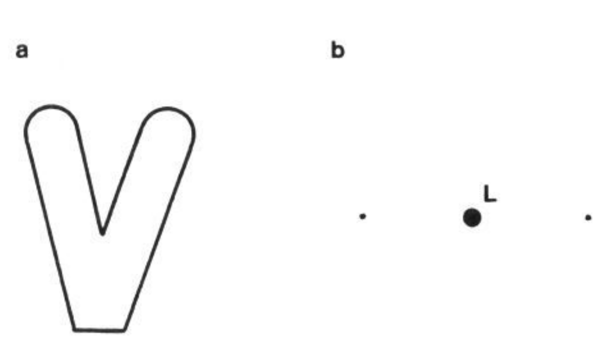

Figure 2 A v-shaped object and its transform

(a) shows a v-shaped object, and (b) shows its transform. The main peak appears at the localisation point L, and the subsidiary peaks appear at equal distances D either side of it: D is equal to the distance between the two prongs of the v.

between the two prongs of the v, and hence there is a local ambiguity for a particular set of pixels about their location within the object. In this case it is not the pixels within the concavity that pose the problem, but those on the two convex regions related by symmetry. This may at first appear paradoxical, but clearly such a situation can only arise when a concavity is present - it is an indirect effect of a concavity that convex regions related by symmetry can arise. On applying our normal rule for accumulating points in parameter space, we find that there are (even in this simple case of 2-fold symmetry) three main peaks in parameter space. The largest is that corresponding to L, at which all points on the boundary of the object have contributed: the other two have equal size corresponding to the much smaller numbers of edge pixels that have contributed to them (Figure 2). A variation on this situation is shown in Figure 3 where the symmetry may be described in terms of two translation vectors in two dimensions.

In such cases, the interpretation strategy that must be applied is to look first for the largest peaks in parameter space, and then to check whether other peaks appear that could possibly have arisen from such symmetries. Since these subsidiary peaks will be at a predictable distance from the main peak, they in fact provide useful additional evidence that the given type of object is present.

It is interesting to consider how detection of such an object will change if it becomes partially occluded. The main possibilities are: (1) the main peak will diminish in size; (2) either or both of the subsidiary peaks will dissappear. If only one of the prongs of the v-shape is visible in the original image, then only one subsidiary peak will be found in parameter space. Now the GHT is known to be highly robust and to have excellent capability for detecting objects that are even quite grossly occluded. In this case, to interpret an image in which only one prong of the v-shape is visible, we must know that it is possible for two small roughly equal peaks to appear in parameter space, either of which may correspond to L. If one of these is larger than the other, it is most likely that it will correspond to L, though some ambiguity will remain. In such cases symmetry has permitted the suppression of information that could have been used to locate the object unambiguously. This phenomenon will clearly generalise to more complex object shapes, where higher order symmetries can be present.

a

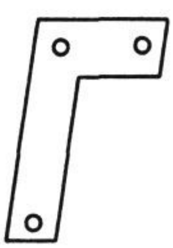

b

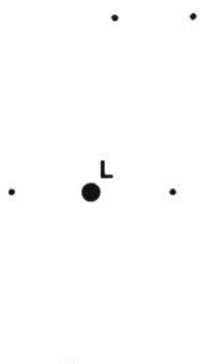

Figure 3 An object containing holes and its transform

In this case the original object (a) is a bracket containing holes, and the symmetry is describable in terms of two translation vectors in two dimensions. (b) shows the transform.

Although symmetry would appear generally to be a disadvantage in locating an object (especially in situations where occlusions can occur), this is not actually so. Consider for example the following case where a periodic structure is being sought (Figure 4). In that case available edge locations provide a barrage of peaks in parameter space, and hence a relatively large amount of ambiguity in the event of occlusions. However, ignoring possible occlusions, there is also an enormous amount of information on the longitudinal location of the object. In fact there is much more information on longitudinal location than there would be for an object of equal size containing no concavities. Furthermore, there is much greater accuracy of longitudinal location than for lateral location, again because of the presence of concavities. We note also that the additional accuracy is not just due to the increase the number of edge pixels for that size of object, but arises partly from the changed distribution of edge orientations within the shape.

V HOW ACCURACY OF OBJECT LOCATION IS AFFECTED BY SHAPE

In this section we model the situation described above by examining the distribution of object boundary orientations. For a circle this will be uniform. It will also be uniform for an object such as a washer which contains a hole and for which this concavity has not changed the capability for measurement of object location - supposing that we normalise to the situation for objects of constant total boundary length. (The reason it is necessary

a

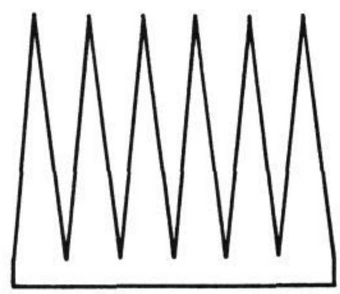

b

Figure 4 A periodic structure and its transform

(a) shows a periodic comb-like structure, which gives rise to a barrage of peaks in parameter space (b). The periodicity provides an enormous amount of information on the longitudinal location of the object, though there is a risk that the transform will be misinterpreted if gross occlusion occurs.
For a linear n-fold symmetry, 2n-1 peaks appear in parameter space; their magnitudes are proportional to the relevant binomial coefficients, $\binom{2n-2}{i}$ $(i=0,1,\ldots,2n-2)$, but (as here) the central peak at L normally has increased size because of the non-periodic parts of the shape.

to normalise our results in this particular way is that on multiplying object linear dimensions by M, we are effectively multiplying the number of independent measurements of x and y object centre location by M, so accuracy is increased, or errors are reduced relatively, by a factor $\surd M$.) For a more instructive example, take the E-shape of Figure 5. In this case the distribution of object boundary orientations θ is $f(\theta)$, which is shown in Figure 6. Taking N_x as the number of vertical boundary points giving information on the x coordinate, and N_y as the number of horizontal boundary points giving information on the y coordinate, we find that $N_y \sim 2N_x$, which means that the y coordinate can be found ~1.4 times more accurately than the x coordinate. Finally, taking the comb shape shown in Figure 4 we obtain the distribution of Figure 7: this has N_x ~$10N_y$, and so the x coordinate can be measured ~3 times more accurately than the y coordinate.

To summarise, we note that object location accuracy depends strongly on object boundary length, and that a preponderance of one or another orientation in the distribution of boundary orientations means that object location will be defined more accurately along one axis than another; in addition, increased accuracy of location parallel to this preferred axis means reduced accuracy of location along a perpendicular axis, for objects of given boundary length.

It will be clear from the above discussion that we can obtain further information on the accuracy of object location by detailed analysis of the boundary orientation distribution $f(\theta)$. Note first that the required information has periodicity π, since an edge at a particular angle α gives the same positional information as one which is inverted - or which has orientation $\alpha + \pi$. If we expand $f(\theta)$ as a fourier series,

$$f(\theta) = a_0/2 + a_1 \cos\theta + a_2 \cos 2\theta + \dots + b_1 \sin\theta + b_2 \sin 2\theta + \dots \qquad (1)$$

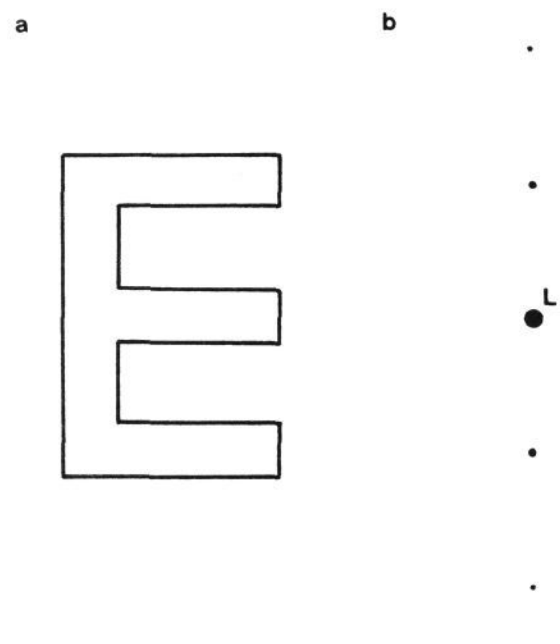

Figure 5 An E-shaped object and its transform

Extended contributions to the transform (b) from each of the straight sides of the original shape (a) are suppressed in this diagram in order to clarify the effects of symmetry.

where

$$a_n = (1/\pi)\int_{-\pi}^{\pi} f(\theta)\cos n\theta \, d\theta \qquad (2)$$

$$b_n = (1/\pi)\int_{-\pi}^{\pi} f(\theta)\sin n\theta \, d\theta \qquad (3)$$

then the above discussion, and the requirement that we should match $f(\theta)$ to a bivariate normal error distribution at L, shows that we are only interested in the variations produced by the terms in $\cos 2\theta$, $\sin 2\theta$. In particular, the term $a_0/2$ gives no variation; the terms in $\cos\theta$, $\sin\theta$ give a variation with the wrong periodicity; and the higher order terms give a variation which is faster than required for the present purpose. Thus a good model of the positional accuracy problem can be obtained by examining how the terms in $\cos 2\theta$, $\sin 2\theta$ modulate the constant term. Particularly important are the amplitude of the variation and the axis along which the variation is largest. By writing the relevant terms in the form

$$g(\theta) = a_0/2 + c_2\cos(2\theta - \gamma) \qquad (4)$$

we find the amplitude as

$$c_2 = (a_2^2 + b_2^2)^{1/2} \qquad (5)$$

the high accuracy axis being given by

$$\gamma = \arctan(b_2/a_2) \qquad (6)$$

Carrying out this computation for a shape in which a proportion p of the boundary has (edge normal) orientation close to α or $\alpha + \pi$ and a proportion $q = 1 - p$ has orientation $\alpha + \pi/2$ or $\alpha + 3\pi/2$, we find:

$$a_0 = (1/\pi)\int_{-\pi}^{\pi} f(\theta)\, d\theta = k(p + q) = k \qquad (7)$$

$$a_2 = (1/\pi)\int_{-\pi}^{\pi} f(\theta)\cos 2\theta \, d\theta = k(p - q)\cos 2\alpha \qquad (8)$$

$$b_2 = (1/\pi)\int_{-\pi}^{\pi} f(\theta)\sin 2\theta \, d\theta = k(p - q)\sin 2\alpha \qquad (9)$$

Hence

$$c_2 = k(p - q) \qquad (10)$$

and $\gamma = 2\alpha$ (11)

If the object has perimeter N pixels, the effective number of pixels providing information about its location parallel to the u-axis (at an angle α to the x-axis) is $N_u = pN$ and the number providing

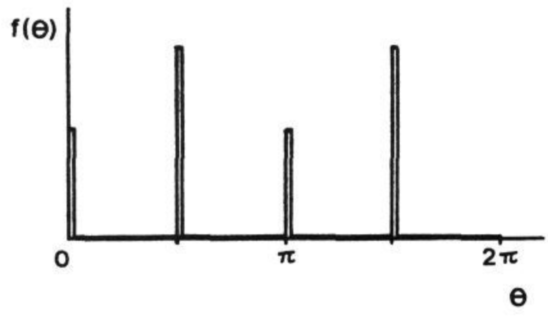

Figure 6 Distribution of boundary orientations for the E-shaped object

information about its location along a perpendicular axis is $N_v = qN$. From equations 7-9 we now deduce

$$p = (a_0 + c_2)/(2a_0) \tag{12}$$

$$q = (a_0 - c_2)/(2a_0) \tag{13}$$

We then get the following values for the effective numbers of pixels:

$$N_u = N(a_0 + c_2)/(2a_0) \tag{14}$$

$$N_v = N(a_0 - c_2)/(2a_0) \tag{15}$$

These formulae clearly give the expected results when $c_2 = 0$ $(q = p)$ and when $c_2 = a_0$ $(q = 0)$, and are valid for any object (such as a line, square, rectangle, box or comb) which has all its edges orientated along or perpendicular to a given axis. However, the model presented here indicates that these formulae are more generally valid for objects of all shapes and orientations. Applying these ideas to some of the features that are commonly used for complex object location and recognition tasks, we find that small holes and corners both have a very small coefficient c_2 and hence are guaranteed to provide reasonable accuracy in all directions.

The fact that we can draw θ-distributions in this way, and deduce measurement accuracies independently of the method of detection, indicates that we have found a fundamental and useful shape descriptor. It also means that our results will have more value than they would have if they were only relevant to some particular specialised technique. On the other hand, the philosophy was obtained from consideration of the GHT, implying perhaps that it is capable of achieving results that are limited more by fundamental properties of shape than by its own characteristics. In this respect we note that the GHT is not an arbitrary technique but is derived from, and closely related to spatial matched filtering, which in turn is known to give optimal signal-to-noise ratio in locating objects (Davies 1987a); in addition, accuracy will be strongly dependent on the possibility of discerning object positions when the reasons for error are fundamentally linked to the presence of noise.

VI THE EFFECTS OF ORIENTATION ERRORS

In spite of what was said in section V about the GHT - namely that its likeness to a spatial matched filter makes it highly efficient at suppressing image noise - there is some doubt about whether it could in practice realise the object location accuracies indicated above. The main reason for this doubt is the fact that every edge pixel is used to provide an independent estimate of the location of L, this estimate being limited by the accuracy with which edge orientation can be measured locally. It has frequently been pointed out that an error $\delta\theta$ in the estimation of edge orientation will at a distance R lead to a positional error $R\delta\theta$ (Ballard 1981, Davies 1987a). This clearly results in an overall lowering of the accuracy of location, though intuition suggests that if L is located in a particular position - e.g. near the centroid of the object boundary (Davies 1987a) - then the overall error will be minimised. It should also be pointed out that edges are frequently 2-3 pixels thick, and some radial error will consequently arise from this source: however, for objects whose linear dimensions are more than ~40 pixels in size, orientation errors in the location of L will predominate - if only because common edge detection operators such as the Sobel permit orientation to be estimated only within ~1°. We shall therefore ignore radial errors and concentrate on azimuthal errors in L in what follows.

To obtain a clearer view of these azimuthal errors, consider Figure 8. In this case we have a circle, and the localisation point is at an arbitrary position near one end of it. Rather than estimating the effects of azimuthal errors directly on L, we here examine the problem from a different point of view. If the estimated orientation at a particular edge pixel is subject to error $\delta\theta$, then there is a local uncertainty about which part of the object boundary the edge pixel is on. In fact we can mark out a section B of the boundary at which the edge pixel could be. (For simplicity we here assume that there is equal probability of the edge pixel being at any of these locations and zero of it being elsewhere on the boundary. This simplyifying assumption helps the calculation, but will not

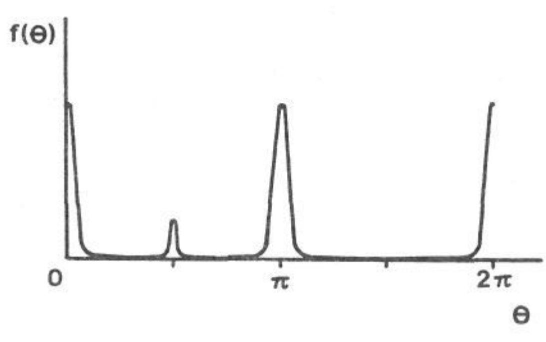

Figure 7 Distribution of boundary orientations for the comb shape

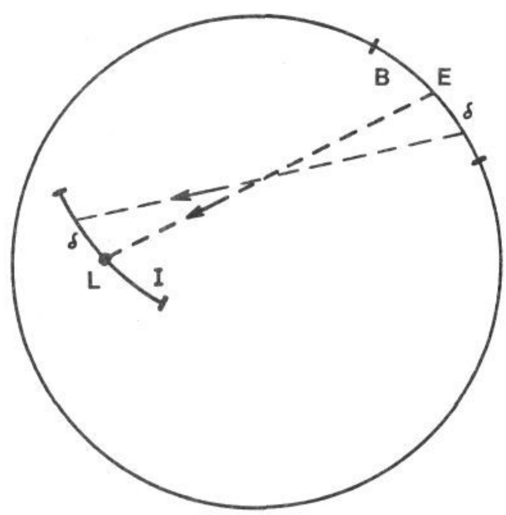

Figure 8 Effect of edge orientation errors on the transform of a circle

Here the transform of a circle is being calculated for localisation point L. Errors in edge orientation at edge point E mean that is not certain which of the points within error arc B it corresponds to. As a result the estimated position of the localisation point L may be any of the points on the inverted error arc I. The final error distribution at L is the result of accumulating all possible inverted arcs of type I.

affect the generality of our conclusions.) Next, we make successive assumptions that the edge pixel is at each of these locations and mark the correspondingly deduced positions of L. For a circle, this means that L must lie on an arc which is an inverted form of the error arc B. Repeating this for all edge pixels gives an isotropic set of error arcs similar to that due to B. It is noteworthy that the fact that some edge pixels are closer to L than others has had no particular effect.

This is a somewhat remarkable result and deserves comment. First, if L is taken as the centre of the circle, the situation reverts naturally to the usual situation, wherein the accuracy can be calculated in terms of the azimuthal error subtended at the centre as a result of the edge orientation errors. Second, we can always assume a best case position for L (here we would take the centre) and then deduce another location for L by adding a constant vector to L <u>or</u> by adding it directly to each of the vectors from the edge pixels to the original L. Thus the reason for the errors being independent of L is that no additional error can result from adding a constant vector to a given vector.

However, it is also necessary to determine why the previous view of the situation was erroneous. Imagine we select a position of L on the circumference of the circle. Then it appears that orientation errors have no affect for nearby edge pixels. But in fact they have a marked effect, since as soon as it is admitted that we do not know exactly what the local edge orientation is we do not know which of a number of edge pixels we are considering. Thus we ought to accumulate a range of possible locations for L in parameter space. This means at least that there is an upper limit on the accuracy with which an edge pixel can predict the position of L. On the other hand there could be a lower limit to the accuracy, which results from edge orientation errors adding to the basic inaccuracy at large distances. However, we have seen, by calculating the azimuthal error as resulting from the ambiguity in the positions on the circumference at which the located edge pixel can be taken to lie, that we obtain a lesser error, which should be regarded as a lower limit on the error of location.

As an example, we consider the situation for an ellipse with semi-axes a and b and high eccentricity. In this case the various arcs of constant angular error on the boundary vary widely in length. If ρ is the local radius of curvature, then the length of the corresponding arc, and of the resulting error arc at L, is $\rho\delta\theta$. This will vary from $(a^2/b)\delta\theta$ to $(b^2/a)\delta\theta$ for various positions of the edge pixel on the ellipse. The result is that the distribution of points in parameter space near L is a very much flattened roughly elliptic shape whose ratio of major to minor axes is $\sim(a/b)^2$. As the original ellipse becomes more elongated the approximation we have taken eventually breaks down, since $(a^2/b)\delta\theta$ becomes greater than 2a. However, it is clear that in that case the longitudinal error becomes equal to 2a - the length of the corresponding straight line - while the lateral error becomes so small that it will be dominated by radial errors due to the width of the object boundary. Thus the model we have adopted gives exactly the right extrapolation to the case of a straight line, as specified by Davies (1987b). On the other hand a calculation based on working out the arc length subtended at L by azimuthal errors resulting from the relevant distance from the edge pixel would predict totally the wrong distribution of errors in L (namely, L would be known more accurately along the length of the line than laterally).

A. <u>Summary and appraisal of GHT positional accuracy</u>

The ideas of this section have so far shown that the position of the localisation point is not a relevant factor in finding the accuracy with which the GHT can estimate object location. They have also shown that edge orientation accuracy is important in helping to determine the position of an object: i.e. the GHT is capable of extracting some edge orientation information to refine the measured position of an object. We now need to work out how these results relate to those of section V.

In fact, section V predicted the relative accuracies with which object location could be found in two perpendicular directions. It did this on the basis that the edge pixels have a given orientation distribution $f(\theta)$. However, we now note that this calculation was based on a model such that a given edge pixel gives a measure of object location normal to the edge, with radial accuracy given by σ_r, but *gives* <u>no</u> information on *object location in the* azimuthal direction. Thus the GHT is well able to meet this requirement and indeed should enable somewhat higher accuracies to be achieved. To this extent the model of section V is inadequate. However, it does permit us to obtain simply a <u>lower limit</u> on the available accuracy of object location. For this we assume a basic accuracy of σ_r, and modify this in two directions taking account respectively of the effective numbers of pixels, N_u and N_v (see section V). Thus the error σ_r will in the two directions be divided by factors $\sqrt{N_u}$ and $\sqrt{N_v}$.

VII <u>CONCLUDING REMARKS</u>

This paper has examined the effects of concavities on the performance of the GHT. Perhaps surprisingly, this seemed to lead to problems only in cases when specific symmetries arose. Then additional peaks in parameter space appeared which could lead to potential ambiguities which might not be resolvable in cases of gross object occlusion. On the other hand the additional peaks could in principle be used to help with object location and identification, and need not be regarded as disadvantageous. Next, it was found that the distribution of edge orientation values permitted the accuracy of object location in different directions to be estimated in a straightforward manner: this approach appeared fundamental and independent of the particular method used to locate objects - thus representing an upper limit on the available accuracy of object location assuming that local edge orientation information is ignored. It was shown that the GHT is well able to meet this accuracy, and should in fact be able to exceed it if local edge orientation information is taken fully into account.

The theory developed here shows how the accuracy with which an object may be located depends on (a) the total number of boundary pixels and (b) their relative orientations. It also makes clear that, for a given boundary length, enhanced accuracy of location parallel to one axis is matched by reduced accuracy of location along a perpendicular axis. In addition, small objects need not be less accurately locatable than large objects if, because of their more complex shapes, they have similar boundary lengths. For example, a small object with many holes or with a sawtooth edge will be very accurately locatable.

In addition, it has been found that the accuracy of location obtainable using the GHT is independent of the position of the object localisation point.

Finding the limits of object location accuracy is vital not just directly - e.g. in getting robots to assemble precision parts - but also for inspection purposes, e.g. when the dimensions of complex machinery must be measured as accurately as possible (within the available resolution) by precise location of individual components and features. In this respect the paper has been able to ascertain a number of the inherent measurement limitations of the GHT and related techniques.

REFERENCES

1 Ballard, D. H. (1981) "Generalising the Hough transform to detect arbitrary shapes", Pattern Recogn., 13, no. 2, 111-122

2 Davies, E. R. (1984) "Design of cost-effective systems for the inspection of certain foodproducts during manufacture", 437-446 in A. Pugh (editor), Proc. 4th Conf. on Robot Vision and Sensory Controls, London (9-11 Oct 1984)

3 Davies, E. R. (1986) "Reduced parameter spaces for polygon detection using the generalised Hough transform", Proc. 8th Int. Conf. on Pattern Recogn., Paris (28-31 Oct 1986), 495-497

4 Davies, E. R. (1987a) "A new framework for analysing the properties of the generalised Hough transform", Pattern Recogn. Lett., 6, no. 1, 1-7

5 Davies, E. R. (1987b) "A new parametrisation of the straight line and its application for the optimal detection of objects with straight edges", Pattern Recogn. Lett., 6, no. 1, 9-14

6 Duda, R. O. and Hart, P. E. (1972) "Use of the Hough transformation to detect lines and curves in pictures", Comm. Assoc. Comput. Mach., 15, no. 1, 11-15

7 Hough, P. V. C. (1962) "Method and means for recognising complex patterns", US Patent 3069654

8 Kimme, C., Ballard, D. and Sklansky, J. (1975) "Finding circles by an array of accumulators", Comm. Assoc. Comput. Mach., 18, no. 2, 120-122

9 Rosenfeld, A. (1969) Picture Processing by Computer. Academic Press, New York