# Application of the generalised Hough transform to corner detection

E.R. Davies, MA, DPhil, CPhys, FInstP

*Indexing terms: Pattern recognition, Image processing*

**Abstract:** A new approach to corner detection is described which is based on the generalised Hough transform. The approach has the advantage that it can be used when objects have curved sides or blunt corners, as frequently happens with food products; in addition, it can be tuned for varying degrees of corner bluntness. The method is inherently sensitive: we have shown how it may be optimised for accuracy in the measurement of object dimensions and orientation.

## 1 Introduction

Many parts of image analysis and industrial inspection require efficient algorithms for locating specific objects in images, with the aims of counting, checking, measuring and generally performing close scrutiny. In industrial applications this can be vital for safety or for quality control, yet profit margins may not be sufficiently great to permit expensive and powerful computational facilities to be used. Indeed, in food product applications the demand can be for hardware costing less than £10 000, while algorithms must be capable of making numerous tests on products during manufacture at rates of 10–30 items per second [1].

In practice, industrial and other images are subject to 'clutter' from objects near the one currently being located, and in any case the background is seldom feature-free. This means that object detection is by no means a trivial operation, and requires both sensitivity and powerful discrimination. Sharp features and well-defined shapes help considerably with object location. Over the past ten years, much attention has been devoted to devising robust algorithms for detecting round objects and circular features such as holes, while straight line detectors and more recently corner detectors have also been the subject of considerable development — see for example References 1 to 4.

This paper is concerned with corner detection. It is relatively straightforward to locate corners of sharply defined metal objects such as nuts, flanges and so on. However, foodproducts and certain other soft or brittle components are more difficult to deal with, as the true corner may be broken off or rounded (Fig. 1). Even with metal objects this can easily occur: indeed, on some occasions the whole purpose of locating objects may be to detect defects of exactly this type. The objective of this paper will be to elucidate the types of algorithm that are required to locate corners which are not guaranteed to be sharp.

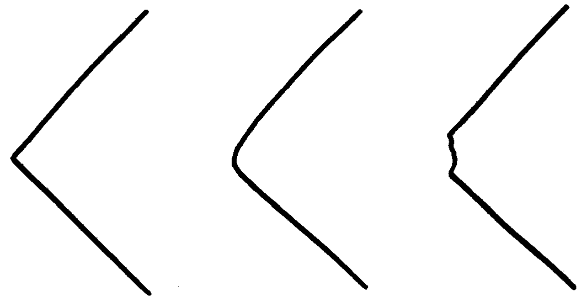

**Fig. 1** *Types of corner*

## 2 Methods for locating blunt corners

As mentioned above, corner location in real industrial images may not be entirely straightforward if there is any possibility that the corners have become blunted by breakage, wear or imprecise manufacture. This situation arises particularly with foodproducts, e.g. when a fish trapezoid is covered with batter or a rectangular biscuit is coated with chocolate or cream: not only may such products have corners different from the norm, but the norm itself may not be very clearly defined.

The obvious way to detect corners in such cases is to idealise the product by locating its sides, and to produce the resulting straight lines until they meet in the presumed corner. For this purpose one requires a straight edge detection algorithm, and the Hough transform forms the basis of most such algorithms [3]. One problem with this approach is that in an image containing several products or objects, many lines will be present, and they will give rise to a rather large number of possible corners, all of which have to be checked (Fig. 2). The reason for this ambiguity is that the basic Hough transform approach does not give longitudinal localisation information on the straight edges it detects. In addition, producing the straight edges is prone to a certain amount of error, especially if there is any curvature in the sides of the object.

For these reasons it is worth seeking a means of detecting the corners directly. In principle this could be

Paper 5743E (E4), first received 12th January and in revised form 24th August 1987

The author is with the Machine Vision Group, Department of Physics, Royal Holloway and Bedford New College, Egham Hill, Egham, Surrey TW20 0EX, United Kingdom

achieved by some sort of template matching technique, but in practice with blunt corners the templates would be so large that the resulting search problem would be intractable (For an $N \times N$ pixel image, and a set of $M$ $n \times n$ pixel templates corresponding to $M$ possible orientations of the corner, some $M \cdot N^2 \cdot n^2$ operations would have to be performed. This would rise to $Q \cdot M \cdot N^2 \cdot n^2$ where there are $Q$ significantly different corner angles.) The problem of heavy computational load must also apply to the 'local ordered grey levels' technique of Paler *et al.* [4] — especially if a large neighbourhood is employed to locate blunt corners. A possible solution seems to be to use the generalised Hough transform (GHT), which has been shown capable of detecting known shapes in images with quite impressive efficiency [5].

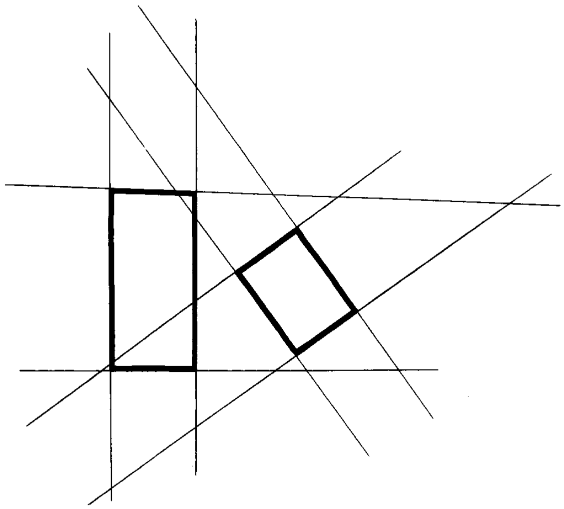

**Fig. 2** *Potential corners that arise from the extended sides of two rectangles*

In the next Section we describe the GHT, and show how it may be used to detect objects with straight edges. Then in Section 4 we apply the technique to corner detection.

## 3 The generalised Hough transform

The Hough transform has shown itself to be extremely valuable for the detection of straight lines and curves such as circles and ellipses [1–3, 6]. In 1981, Ballard generalised the method for the efficient detection of arbitrary shapes [5]. Recently, Davies has found certain problems with the GHT, and has shown how to optimise its sensitivity, e.g. with regard to point weighting in parameter space [7]. He has also shown how objects having straight edges may be detected optimally using the GHT schema [8] (see below).

Before describing the action of the GHT it is necessary to define the term 'localisation point': this is merely a convenient reference point within an object which may be used to help locate it.

The first stage in computing the GHT of an object is use of an edge detector to find the pixels around the outside of an object. For each edge pixel, the orientation is estimated, and the position of a localisation point $L$ is computed by moving a distance $R$ from the edge pixel in a direction $\phi$ which is a function of the direction $\theta$ of the edge normal (see Fig. 3). This operation is repeated for every edge pixel, and a set of candidate locations (or 'votes') is found which is then averaged to determine the position of $L$ as accurately as possible.

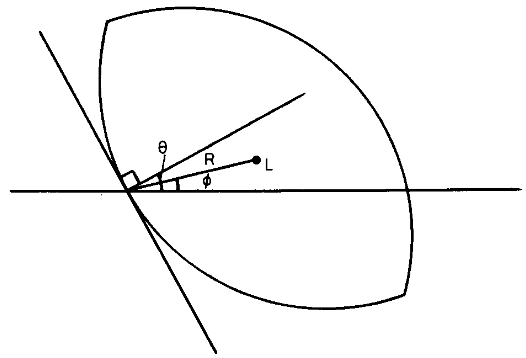


**Fig. 3** *Computing the generalised Hough transform*

Both the distance $R$ and the bearing $\phi$ of $L$ relative to the edge pixel are functions of the edge normal orientation $\theta$. In the GHT, values of these functions are obtained analytically or from a lookup table, depending on how the object shape is defined. The result is that the method can be used to locate an arbitrary shape. Many shapes, such as ellipses and parabolas, can be expressed analytically, but the majority cannot. The great step forward made by Ballard was to formulate the GHT rather generally, so that it could be applied to detect an enormous number of industrial and other objects.

### 3.1 Problem with straight edges

It turns out that this form of the GHT is suboptimal for detecting objects possessing straight edges [8]. We can see this, for example, when attempting to locate a semicircle of known orientation, using a localisation point $L$ near the centre of the object (see Fig. 4). Every edge pixel on the circular part of the perimeter contributes to the peak in parameter space*, at the required location $L$. However, only one point on the straight part of the periphery contributes to this peak: all other pixels on the straight portion contribute to a set of minor peaks and do not help in locating the object. For a square the situ-

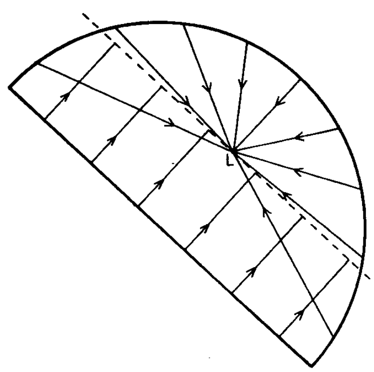


**Fig. 4** *Locating a semicircle using the GHT*
Here the arrows indicate the directions in which the localisation point L is assumed to lie for the typical edge points

* It is conventional to call the transform space 'parameter space', even when it is, as here, congruent to image space.

ation is even worse, since only about four edge pixels are likely to contribute plots to the main peak instead of (typically) well over a hundred.

Davies [8] showed that to guarantee achieving optimal sensitivity in such cases, each edge pixel on a straight edge *must* be transformed into a line of points equal in length to the whole length of the line in the original shape, and of identical direction. This procedure involves a certain amount of additional computation, but the overall algorithm is still efficient. The resulting transforms in parameter space frequently look rather unusual (see Fig. 5), but the method has the advantage of giving provably optimal sensitivity for detecting an object of the selected shape, since it is formally equivalent to using a spatial matched filter [7–9]. A further advantage of this approach to line detection is that line transforms possess (and represent) a restricted longitudinal localisation, unlike the situation for the normal Hough transform line detection schemes (see Section 2).

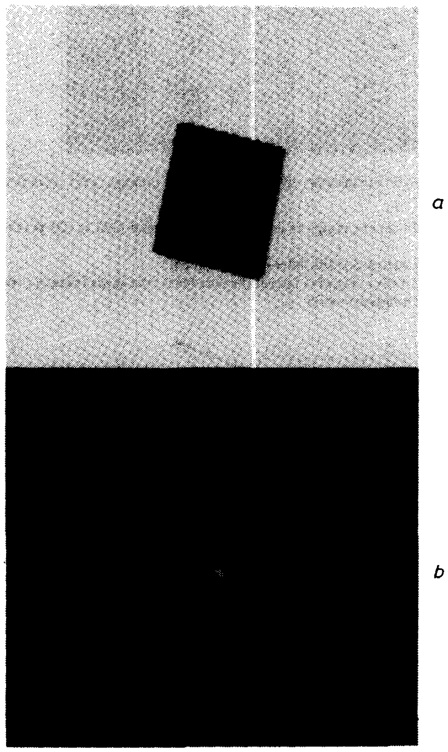


**Fig. 5** *The transform applied to a square*

*a* Original pattern: off-camera image of a high contrast square object (128 × 128 pixels: 64 grey levels)
*b* Transform in parameter space
Here the centre of the square is chosen as the localisation point *L*; and each edge point in (*a*) is transformed into a line of points parallel to the edge passing through *L*. Thus all edge points on the square contribute to the main peak at *L* in parameter space (*b*)

## 4 Use of the generalised Hough transform for corner detection

In order to apply the GHT to corner detection it is possible to use the GHT in its usual line detection mode, as for square object detection, and to look for peaks in parameter space. Two lines then appear in parameter space for each corner, and the peak occurs at the crossover. Since parameter space is congruent to image space, the peak is at the position of the corner — or if the corner is blunted, at the position of the idealised corner.

Unfortunately, this simple approach to corner detection using the GHT is not as accurate as might be wished. Essentially this is because it involves extrapolating to the position of the corner: if interpolation were possible instead of extrapolation, this would be a significant improvement. With this in mind we have modified the method by selecting a corner localisation point that is *inside* the object. This is achieved very simply — by introducing a lateral displacement of the line transform of each edge point, so that most of the transform falls within the object boundary (see Fig. 6). Then the peak at the crossover is formed by interpolation.

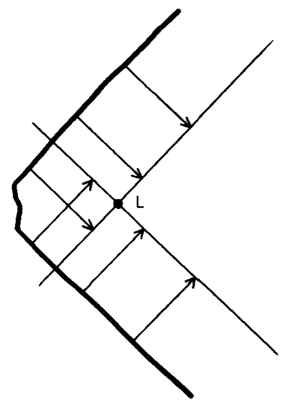


**Fig. 6** *Finding corners by an interpolation scheme*

Here the arrows indicate the lateral displacement of the line transforms: the peak at the crossover point is the localisation point L of the corner. In effect, L is located by a process of interpolation: as a result it is offset from the position of the idealised corner

To select a suitable lateral displacement, consider the situation for a square. If the displacement is zero, the corner localisation points are situated at the corners of the square. If the displacement is equal to half the side of the square, then the localisation points are all at the centre of the square. Clearly, some intermediate lateral displacement is required (see Fig. 7). A simple solution is to use a displacement equal to the maximum distance over which the corners of the square could be blunted. In practice, slightly greater displacements are found to give greater sensitivity (see Section 5).

The resulting positions of the peaks in parameter space evidently provide accurate information on the positions of the idealised corners of the object (here a square), and there are several ways in which the actual positions of the corners can be recovered. One is to use the structure around the peak to indicate in which of four directions the corner could be situated, and to move an appropriate distance (related to the lateral displacement already used) in these directions; then to check each location in turn to find which corresponds to the true corner. Another is to work out from the set of peaks in parameter space where the object must be, and then (but only if necessary) to deduce the actual positions of the corners. This second approach is the more sound one, and is the better one to use to detect the position and orientation of known objects. For unknown objects there is the possibility of working simultaneously in two parameter spaces, one as indicated above and the other one employing a smaller lateral displacement. Then true corner positions can quickly be determined.

Finally, it should be noted that we have implicitly assumed above that objects are convex and hence that all corners are convex. Clearly, if some concave corners exist, then they have to be searched for using a negative lateral displacement. It is impracticable to use both positive and negative lateral displacements in the same parameter space, since this will lead to a great proliferation of false lines and peaks. Instead two parameter spaces have to be used, one indicating the presence and location of convex corners, and the other indicating the positions of concave corners. We have found that this technique works adequately in practice (see Fig. 8).

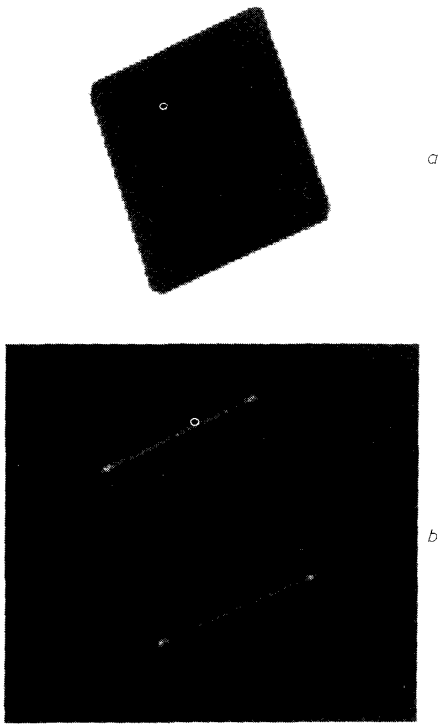

**Fig. 7** *Detection of blunted corners*

*a* Original pattern: off-camera image of a square object with blunted corners (128 × 128 pixels: 64 grey levels)
*b* Transform obtained by using lateral displacements ≃ 10% of the linear dimensions

## 5 Discussion: sensitivity and accuracy

We here consider in more detail the choice of lateral displacement in the corner transformation. There are several disadvantages in setting the lateral displacement to zero: first, it gives significantly reduced sensitivity (in a signal-to-noise ratio sense) in locating the corner; secondly, as we have seen, it gives reduced accuracy of corner location and thirdly, it results in lack of longitudinal localisation of the lines in parameter space, and hence may give rise to spurious corner peaks (see Fig. 2). The first disadvantage arises since the transforms of fewer edge points reach as far as the localisation point: this disadvantage can be eliminated by extending the line transform of each edge point, but this reduces the longitudinal localisation of the lines in parameter space further, and results in a proliferation of spurious corner peaks (see Fig. 2). Thus the first and third disadvantages are linked, and are both reduced progressively by increasing the lateral displacement in parameter space from zero. As indicated in Section 4, the second disadvantage is similarly improved by increasing the lateral displacement from zero. Fig. 9 shows in more detail the effects of inaccuracies caused by errors in estimating edge orientation. With zero lateral displacement, the corner peak is found by extrapolation of the edge direction, and the errors are directly proportional to the angular error; with moderate lateral displacements, errors in estimation of the edge orientation cancel out in first order, since the lines in parameter space are essentially envelopes rather than extrapolations (see Fig. 9).

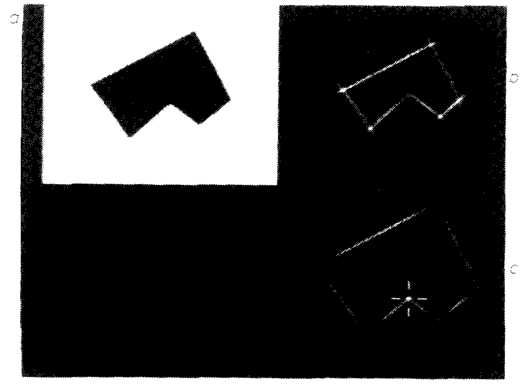

**Fig. 8** *Use of two parameter spaces to detect convex and concave corners*

*a* Original pattern: off-camera image of a flat polygonal object (128 × 128 pixels: 64 grey levels)
*b* Transform obtained with a positive lateral displacement
*c* Transform obtained with a negative lateral displacement: the cross indicates the (offset) position of the concave corner

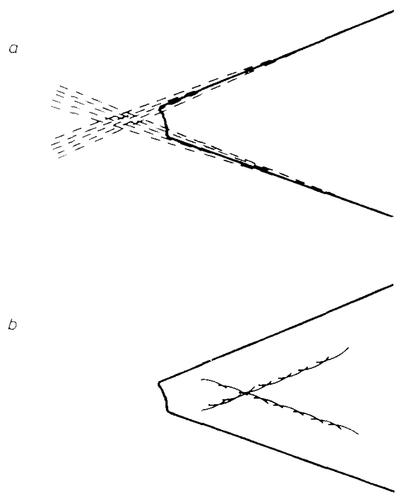

**Fig. 9** *Effect of edge orientation errors*

*a* Case of zero lateral displacement
*b* Case of moderate lateral displacement

The above arguments imply that optimum sensitivity, accuracy and localisation of the corner transform occur when the localisation point is at the centre of the object. Obviously this is not the whole story, since a major purpose of locating the corners of an object rather than its centre is to find as accurately as possible what its orientation is: another purpose is to permit its dimensions to be measured. From these points of view, maximising lateral displacement is counterproductive. Evidently a compromise is required, so that for the types of blunt corner that arise in practice, object size and orientation measurements can be as accurate as possible. Intuition suggests that using about one quarter of an object side for corner detection will give a suitable balance between sensitivity of object location and object orientation. This is backed up by theoretical analysis in Section 5.1. The tradeoff is also affected by any curvature of the lines in the vicinity of the corner. The whole problem then becomes one in which an optimum solution has to be sought using intuitive judgements of relevant corner parameters. We can at least list a sequence of steps which will form guidelines for tailoring the method to individual applications. The two algorithm parameters that are involved are lateral displacement and the length of the line that each edge pixel transforms into. Thus a valid sequence of steps is the following

(*a*) determine the expected bluntness of the corners

(*b*) determine the expected curvatures of the sides in the vicinity of the corners

(*c*) determine the range of corner angles that are expected

(*d*) specify the portions of the sides from which corner positions can best be estimated

(*e*) specify the closest positions that corners of neighbouring objects might occupy

(*f*) deduce the best compromise value for the lateral displacement $D$

(*g*) deduce the best compromise length $T$ for the edge pixel transform

We return in Section 5.1 to means of optimising the choice of the algorithm parameters $D$ and $T$. Meanwhile we illustrate the effect of varying $D$ for a particular type of biscuit (Fig. 10), and we briefly discuss two other points about the method.

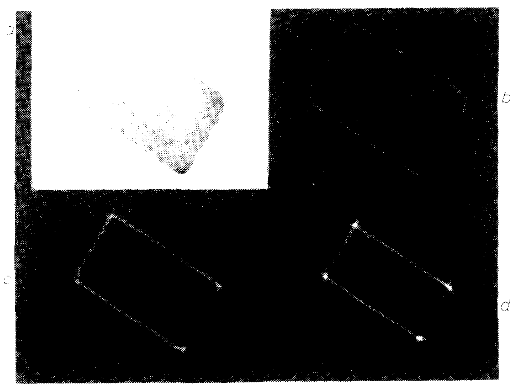


**Fig. 10** *The effect of varying lateral displacement in practice*
*a* Original image of a biscuit (128 × 128 pixels: 64 grey levels)
*b* Transform with zero lateral displacement
*c* Transform with lateral displacement ≃ 11% of the shorter side
*d* Transform with lateral displacement ≃ 22% of the shorter side
Here (*d*) is close to the optimum for this type of object: cases (*b*) and (*c*) are only shown for comparison

First, the technique can be extended if the expected range of corner angles is excessive and is liable to reduce the effectiveness of the method significantly. In that circumstance, several parameter spaces can be used, each with its own value of lateral displacement and transform length. Our research indicates that more than two such spaces are seldom needed.

Secondly, it should be remarked that it is possible to give lines a longitudinal displacement, so that the corner can be detected in its actual position with optimum sensitivity (i.e. optimum signal-to-noise ratio). However, this approach does not eliminate the second and third disadvantages. In addition, it requires that assumptions be made of the corner orientation, and a great many parameter spaces each corresponding to one orientation of the object will have to be constructed and searched through for corner peaks. We have concluded that this is not a viable option and have ignored it for practical applications.

### *5.1 Model for optimising sensitivity and accuracy*

The above discussion of sensitivity and accuracy was carried out on a rather intuitive basis. Here we present a simple mathematical model for optimising corner detection.

Fig. 11 shows an object with blunt corners, the region of bluntness being within a distance $B$ of the idealised

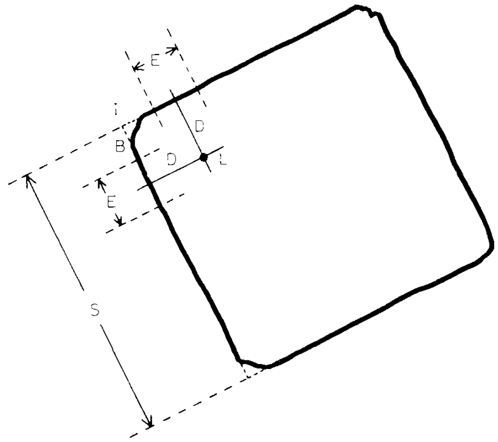


**Fig. 11** *Application of the algorithm to an object with blunt corners*
*I* idealised corner point
*B* region of bluntness
*E* distance over which position of corner is estimated
*D* lateral displacement employed in edge pixel transform
*L* localisation point of a typical corner
*S* length of a typical side
The degree of corner bluntness shown here is common amongst foodproducts such as fishcakes

corner position. Suppose that we wish to make our estimates of corner position over a further distance $E$ along each of the adjoining sides (Fig. 11). Next, assume that the transform of any edge pixel is a line of length $T$ parallel to the local edge orientation and that (as recommended earlier) longitudinal displacements are not to be employed. The only other relevant parameter is the lateral displacement $D$.

Since we only wish edge pixels over distances $E$ to contribute* to the peak at the corner localisation point $L$, by reference to Section 3.1 we specify that

$$T = E \tag{1}$$

In addition $D$, $B$ and $E$ are simply related by geometry (see Fig. 11):

$$D = B + E/2 \tag{2}$$

so it is clear how the algorithm parameters D and T depend on the picture data.

It now remains to calculate the sensitivity and accuracy. The sensitivity will depend closely on the height

* Other edge pixels are not merely irrelevant, but can be regarded as adding noise to the desired signal. This is particularly important if the sides of the object are at all curved.

$H$ of the peak at the corner localisation point $L$, and we know that $H$ is proportional to $E$. We now take the 'signal' as $H$ and to find the 'noise' note that in the region of the peak in parameter space it arises as statistical variations in the numbers of votes accumulated locally in parameter space: hence it is reasonable to assume that the noise is proportional to $\sqrt{H}$. Thus the signal-to-noise ratio $\sigma$ (or 'sensitivity') is proportional to $\sqrt{H}$, and can be expressed as

$$\sigma = \sigma_0 \sqrt{E} \tag{3}$$

Next, the accuracy of corner location will depend on the signal-to-noise ratio $\sigma$. We assume it is proportional to $\sigma$. (There are strong reasons supporting this, as we are effectively averaging our measurement of the relevant part of each side over $E$ independent signals.) However, absolute location of corners may be less relevant than finding the dimensions of the object and its orientation. These are both proportional to the distance between corner localisation points as well as being affected by signal-to-noise ratio. Thus the overall accuracy for measurement of object length and orientation as determined by corner location is proportional to

$$A = (S - 2D)\sigma = (S - 2D)\sigma_0 \sqrt{E} \tag{4}$$

where $S$ is the length of a typical side. Eliminating $D$ now gives

$$A = (S - 2B - E)\sigma_0 \sqrt{E} \tag{5}$$

Differentiating, we find

$$\begin{aligned} dA/dE &= (S - 2B - E)\sigma_0 \sqrt{E}) - \sigma_0 \sqrt{E} \\ &= (S - 2B - 3E)\sigma_0/(2\sqrt{E}) \end{aligned} \tag{6}$$

We can optimise accuracy by making $dA/dE = 0$, so that

$$E = S/3 - 2B/3 \tag{7}$$

This makes

$$D = S/6 + 2B/3 \tag{8}$$

If $B = 0$ this gives $D = S/6$, and for $B = S/8$ we find that $D = S/4$. This is in line with the result we obtained earlier, and is also backed up by our experience with real image data.

Finally, we note that there is no similar functional optimum of sensitivity, since it apparently increases without bound (eqn. 3). However, the sides of the object are limited in length, and hence optimum sensitivity is given by $E \simeq S$ rather than $S/3$, and by $D \simeq S/2$ rather than $S/6$. In fact, optimising sensitivity is seen to give the same condition as that for whole object detection using the GHT: on the other hand, we should remember that a major reason for employing corner detection is to obtain accurate estimates for object dimensions and orientation — even if this implies a loss (by a factor $\sim\sqrt{3}$) in the overall sensitivity of object detection.

In this section we have assumed that errors arising from inaccuracies in edge orientation are small. As such errors increase it will be necessary to reduce $T$ and $E$ progressively, and accuracy will fall. In the limit of increasing orientation errors, we conclude that $D$ will need to be close to $B$. However, such extreme cases do not seem to arise in practice (see for example Fig. 10).

## 6 Conclusion

This paper has described a new method for detecting corners of objects: the method is particularly well adapted for cases where the corners are liable to be blunted for one reason or another. It has been found that the method can be tuned for various degrees of bluntness, and also for different ranges of corner angle. In addition, the method can be adapted for cases where the sides of objects are curved, or where images may contain several objects in close proximity.

The method is based on the generalised Hough transform, which is known to give optimal sensitivity in the location of objects. For corner detection, sensitivity is necessarily reduced, but this is because the algorithm is optimised for accuracy of measurement of object dimensions and orientation, rather than for sensitivity of whole object detection. This reflects the fact that the purpose of corner detection is at least as much for measuring the dimensions and orientation of an object as for locating it.

Setting up of the algorithm is carried out in a prescribed sequence of steps which tailor the algorithm to the application. Though intuitive judgements may have to be made about the values of relevant corner parameters, formulae have been obtained which permit optimal values of algorithm parameters to be chosen systematically.

When compared with other well-known corner detectors such as [4], the method presented here offers the advantage of locating corners in well-defined positions which are invariant to moderate changes in corner bluntness. Overall, the method is sensitive, robust, accurate and efficient: it is also suitable for the detection of corners having a wide range of angles and orientations.


## Acknowledgment

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


## 8 References


1 DAVIES, E.R.: 'Design of cost-effective systems for the inspection of certain foodproducts during manufacture' in PUGH, A. (ed.), Proc. 4th Conf. on Robot Vision and Sensory Controls, London (9–11 Oct 1984) pp. 437–446

2 KIMME, C., BALLARD, D., and SKLANSKY, J.: 'Finding circles by an array of accumulators', *Comm. Assoc. Comput. Mach.*, 1975, **18**, pp. 120–122

3 HOUGH, P.V.C.: 'Method and means for recognising complex patterns', 1962, US patent 3069654

4 PALER, K., FÖGLEIN, J., ILLINGWORTH, J., and KITTLER, J.: 'Local ordered grey levels as an aid to corner detection', *Pattern Recogn.*, 1984, **17**, pp. 535–543

5 BALLARD, D.H.: 'Generalising the Hough transform to detect arbitrary shapes', *Pattern Recogn.*, 1981, **13**, pp. 111–122

6 TSUJI, S., and MATSUMOTO, F.: 'Detection of ellipses by a modified Hough transform', *IEEE Trans.*, 1978, **C-27**, pp. 777–781

7 DAVIES, E.R.: 'A new framework for analysing the properties of the generalised Hough transform', *Pattern Recogn. Lett.*, 1987, **6**, pp. 1–7

8 DAVIES, E.R.: 'A new parametrisation of the straight line and its application for the optimal detection of objects with straight edges', *Pattern Recogn. Lett.*, 1987, **6**, pp. 9–14

9 SKLANSKY, J.: 'On the Hough technique for curve detection', *IEEE Trans.*, 1978, **C-27**, pp. 923–926


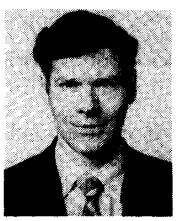


**Dr. E.R. Davies** graduated from Oxford University in 1963 with a first class honours degree in physics. After twelve years' research in solid state physics, he became interested in vision and now leads the Machine Vision Group at Royal Holloway and Bedford New College (London University), where he is a senior lecturer. He has worked on thinning, edge detection, Hough transforms and robust pattern matching methods. Much of his recent work has been on automated visual inspection, particularly for the food industry.

Dr. Davies is a member of the British Pattern Recognition Association, the British Robot Assocation and the IEEE Computer Society. He is a Fellow of the Institute of Physics.