Gradient based progressive probabilistic Hough transform

C.Galambos, J.Kittler and J.Matas

Abstract: The authors look at the benefits of exploiting gradient information to enhance the progressive probabilistic Hough transform (PPHT). It is shown that using the angle information in controlling the voting process and in assigning pixels to a line, the PPHT performance can be significantly improved. The performance gains are assessed in terms of repeatability of results, a measure that has direct relevance for its use in many applications. The overall improvement in output quality is shown to be greater than that found for the standard Hough transform using the same information. The improved algorithm gives results very close to those of the standard Hough transform, but requires significantly less computation.

1 Introduction

The Hough transform is a well studied method for extracting geometric primitives. In this paper we look at some refinements to a variant of this method called the 'progressive probabilistic Hough transform' (PPHT). Details of this method have been presented previously [1–3].

The key difference between the PPHT and the probabilistic Hough transform (PHT) [4] on which it is based is that the accumulator space is scanned for significant peaks as each vote is cast and lines are removed as they are found. When a line is detected, all edgels that are assigned to the line are removed from the list of unused edgels. Votes of the edgels that have been added to the accumulator are removed. This scheme compares well in terms of computational cost with other adaptive schemes [5–7] which rely on monitoring peaks in the accumulator space. It also avoids making assumptions about the distribution of peak sizes, and problems that may occur if these constraints are violated.

The only parameter the PPHT requires to be set is the false positive threshold (*l*). The threshold controls the fraction of acceptable false positives that can be generated by the algorithm. The effect of this parameter on the total number of votes required to process an image is studied in [2]. In essence, the smaller the value, the longer it will take for the algorithm to terminate but the false-positive line detection will also be lower.

In this paper, we investigate the benefits gained by using gradient direction in conjunction with the PPHT algorithm. Similar improvements are possible with the SHT [8, 9] and as they have been considered elsewhere [10–12], their

performance benefits will not be examined in detail here. We show that with the use of gradient direction information, the accuracy of the PPHT can be brought very close to that of the SHT using the same information, while retaining a significant advantage in computational speed.

2 Modifications to the PPHT algorithm

The use of gradient direction does not significantly alter the overall operation of the PPHT algorithm (algorithm 1). The differences will be explained later in this section.

Algorithm 1: Progressive probabilistic Hough transform

Step 1. Vote into the accumulator with a single edgel randomly selected from the input set. If none, then finish.
Step 2. Remove edgel from input set.
Step 3. Check if the highest bin count in the accumulator (the peak) that was modified by the new edgel is higher than threshold *thr(N)*, defined in Section 2.1. If not, then go to 1.
Step 4. Look along a corridor defined by the peak, and find the longest segment of edges that is either contiguous or not exhibiting a gap larger than a given threshold. This forms the segment support.
Step 5. Remove the edges in the segment support from input set.
Step 6. Unvote from the accumulator all the edgels in the segment support.
Step 7. If the segment support is longer than the minimum length add it into the output list.
Step 8. Goto 1.

In using gradient direction, we make two key changes to the original PPHT algorithm. The first is to constrain the range of angles, which has two effects. It reduces the computation required to process a new edgel; it also reduces the clutter in the Hough space, or, in other words, increases the signal to noise ratio. Gradient direction can be made available from several sources. Most edge detectors provide the information as part of their output. If edge information is not available directly, it is possible to estimate it, for instance by calculating the moments of the neighbouring edgels.

The second change is to use gradient direction in the post-processing part of the PPHT. When the corridor corresponding to a peak in the accumulator is searched, only edgels with a gradient consistent with the line parameters specified by that peak (within the same range used for voting) are considered for assignment to the line. This can significantly improve the quality of lines generated, both in edgel assignment at T-junctions, and where there are many closely spaced lines.

The key factor influencing the improvement gained by using gradient direction is the amount of uncertainty of the gradient direction at each edgel. Unfortunately this can be very difficult to quantify theoretically for the whole image capture and processing chain. Where it is possible, the results would still only be applicable to a single set of equipment. For the purposes of this paper, we will assume the angular error distribution is unknown but approximately uniform for each edgel.

2.1 Changes to the threshold computation

Constraining the number of angles used when voting changes the distribution of votes in the accumulator, and so the original analysis of the peak detection threshold for PPHT is no long accurate. It is now easier to consider the accumulator to be a set of 1-D histograms of possible ρ values, one histogram for each angle θ considered (the standard ρ–θ representation of the accumulator space [13] is assumed). When a vote is cast only a subset of these histograms is updated.

In setting the decision threshold, we assume that all points are due to noise. It is a worst-case assumption, but if many lines are present in the image the assumption is almost valid, since only a fraction of points belong to any single line.

Let us denote the number of histograms for angles as N_θ the number of bins for distance from origin as N_ρ and the number votes for a given angle as V_θ. We adopt the following model of the voting process. Every randomly selected edgel votes once into a subset of the available N_θ histograms, the exact bins depending on the gradient direction of the edgel and the uncertainty in its value. The size of the uncertainty, denoted by γ, is related to V_θ as follows:

$$\gamma = \frac{360 V_\theta}{N_\theta} \qquad (1)$$

An edgel can belong either to no line (a noise point), to a single line, or lie on an intersection of lines. In the first case all votes cast add noise to those histograms. We assume that, for every one of the V_θ, histograms voted into, a random bin is incremented. Each bin in a histogram is equally likely to be incremented with probability $1/N_\rho$. If a point lies on a line, one vote is cast into the bin corresponding to this line and the remaining V_θ, votes are assumed to fall into random bins, one in each of the corresponding histograms. For points on line intersections, we assume $V_\theta - 2$ votes fall in random bins. Since $V_\theta \gg 1$, $V_\theta \approx V_\theta - 1 \approx V_\theta - 2$, we do not (and we cannot) distinguish between the three cases and assume that always random V_θ bins are incremented.

Clearly, the ρ histograms are not independent, and the counts in bins with similar θ and ρ are not statistically independent either, because of the cosine voting pattern of a single point. Nevertheless, to keep computation tractable, we will assume that the count in any single bin in a histogram is an independent random variable with binomial distribution B (V_θ, p), where V_θ is the number of

edgels that voted in each ρ histogram so far and $p = 1/N_\rho$ is the probability of selecting a particular bin with a given θ. In our voting model, the distribution of votes in the N_ρ bins for a given θ follows the multidimensional hypergeometric distribution (not multinomial distribution, since the sampling is without replacement). We adopted the $B(V_\theta, p)$ simplification because we could not find a practical (efficient) testing procedure for the hyper-geometric distribution. The hypothesis that is being tested after every bin increment is the following:

Is the count $C(\rho, \theta)$ in bin (ρ, θ) higher than a value likely to occur if $C(\rho, \theta)$ was a realisation of a random variable with binomial distribution $B(V_\theta, p)$?

We would like to set the threshold so that

$$P(C(\rho, \theta) > thr(V_\theta)) < l \qquad (2)$$

Significance level l is a user parameter that shall, if the model is accurate, indicate the number of false positives (in case of no post-processing) due to noise. If there is more than one peak found in voting for a single edgel, the highest is used. For a binomial distribution, it is easy to compute the threshold for a given N by evaluating the sum

$$\sum_{i=0}^{j} P(C(\rho, \theta) = i) \qquad (3)$$

for all j till $1 - l$ is reached. This value can be computed relatively efficiently with the incomplete beta function [14]. To further speed the process of voting, these values can be precomputed and stored in a look-up table.

3 Experiments on synthetic data

The experiments presented aim to quantify the performance benefits gained by using gradient direction with the PPHT. The main factors for consideration are the improvement in the interpretation of the results and computational benefits.

For consistency, the stopping criterion of the algorithm exploiting gradient direction should be modified to take account of the reduced influence of noise on the accumulator statistics. However, to make clear performance comparisons between the modified and unmodified PPHT, no changes were made to the threshold calculation in the first set of experiments.

The synthetic images used for these experiments were 256 pixels squared, each with 20 lines of random length uniformly distributed between 1 and 100 edgels. An example of such a synthetic image can be seen in Fig. 1. Each experiment was repeated 100 times. This enables the computation of the means and the standard deviations for the measured quantities shown in the graphs. All error bars correspond to one standard deviation. To make the gradient direction data in the synthetic image realistic, it was estimated by counting the moments of all the edgels within a radius of 2.5 grid squares. This estimation also works well on real images, and sometimes better as they often contain fewer crossing lines.

The following criteria were used for determining the error statistics. False positives are all those lines detected that cover less than 80% of any single ground-truth line in the image. False negatives are those lines in the model that are covered by less than 80% by the detected lines, excluding those counted as false positives.

Fig. 2 shows the number of voting operations used in processing the images. Though voting for a restricted range
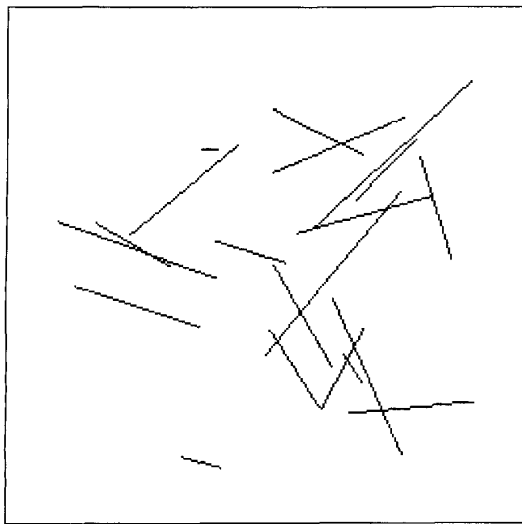
Fig. 1 *Example of synthetic edge image*

Fig. 3 *SHT with gradient direction*
——— false negatives
– – – false positives

of angles means that the number of voting operations is no longer directly proportional to the time required for computation, these numbers can still be used to compare the relative performance of the algorithm. The test images contain 2000 edgels, and hence the full SHT uses 2000 voting operations. This means the results for the PPHT were obtained with about one-tenth of the operations required by the SHT.

As the range of angles used drops below the uncertainty in the gradient direction information, the number of votes needed to process the image starts to rise for values of γ below 30° (Fig. 3). This occurs because, as the range of bins incremented becomes smaller than the uncertainty in the direction, it becomes increasingly likely that the bin corresponding to the actual line parameters will not be incremented.

Figs. 4 and 5 show the detection performance results for the PPHT as a function of the gradient angle constraint. For
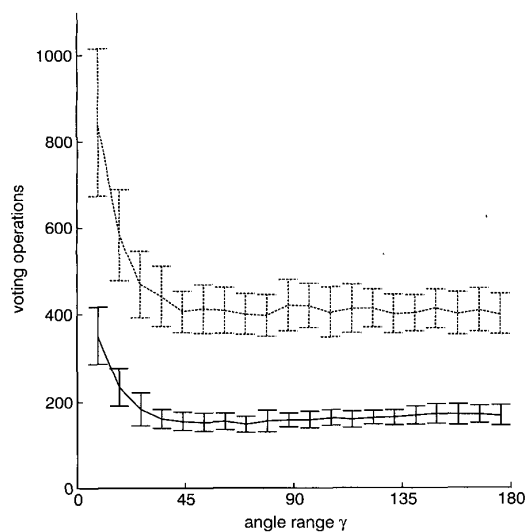
comparison, the results of SHT using the same information are shown in Fig. 3. The performance is measured in terms of average number of false positives (Fig. 4) and false negatives (Fig. 5) (undetected lines) as compared with the known ground truth for each test image. We note that when γ has values between 30 and 60°, the number of false negatives dips significantly for the faster version of the PPHT with the false positives threshold l set at 0.1. At the same time, the false positive rate is significantly reduced to a level comparable to the PPHT operating at the high l of 10^{-9}. It is important not to set the orientation threshold too tight, as the false negative rate dramatically increases as the angle uncertainty interval approaches zero. Fortunately, the performance curves are reasonably flat for values of γ between 30 and 60°, and one can allow a sufficient margin to prevent moving into the degraded performance range due to changes in the image signal to noise ratio.

Fig. 2 *PPHT number of edgels that voted, using gradient direction for accumulation*
——— $l = 0.1$
– – – $l = 10^{-9}$

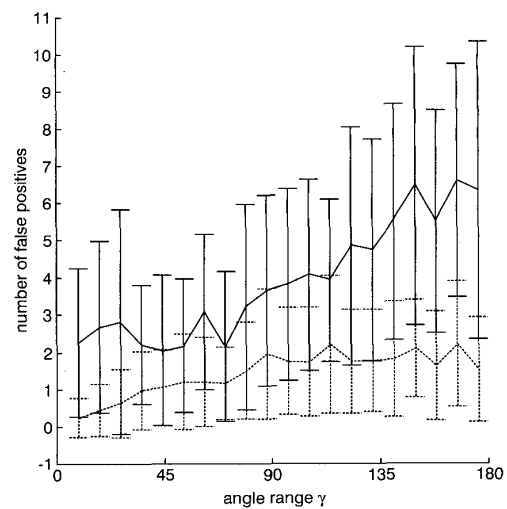**Fig. 4** *PPHT false positives, using gradient direction info for accumulation*
——— $l = 0.1$
– – – $l = 10^{-9}$

160
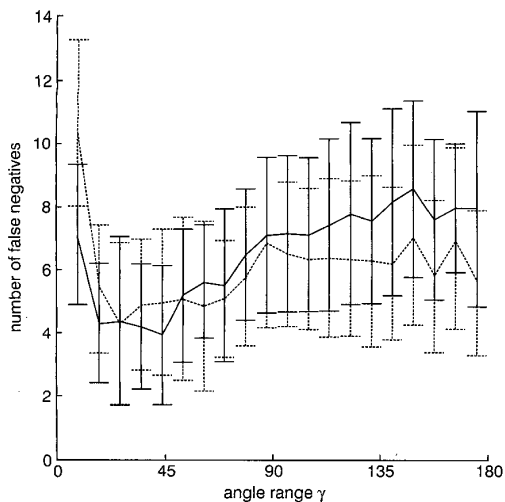
Fig. 5 *PPHT false negatives, using gradient direction info for accumulation*

—— $l = 0.1$
‐‐‐ $l = 10^{-9}$

Fig. 6 *Comparison of results from PPHT and the SHT on real images using gradient direction*

—— false negatives
‐‐‐ false positives

For values of l smaller than 10^{-4}, the results for false positives shown in Fig. 4 are fairly similar for the PPHT and the SHT, but the number of false negatives in Fig. 5 shows about a 20% drop over the previous results for values of γ between 30 and 60°. For small values of l the overall results for the PPHT, at least on synthetic images, are better than those for the SHT. As can be seen from Fig. 2, this extension has little impact on the number of voting operations required to process the image.

4 Experiments on real images

In the first experiment on real images, we compare the output of the SHT and the PPHT on a real image. This experiment illustrates that the improvements give similar benefits when processing real image data. These experiments were run with the house edge image as used in [9]. It is worth noting that the PPHT with gradient direction information has been tested successfully with many other real images.

The SHT was run with γ of 30° to give as near optimal interpretation as possible. The SHT and the PPHT use different stopping rules and hence the number of short lines recovered vary significantly. To reduce problems with this causing excessive false positives, only lines of 10 edgels and longer where used in the comparison. The PPHT was run with an l of 10^{-4} which has been found to give good performance.

Fig. 6 shows the results of the comparison. As indicated by the experiments on synthetic data there is an optimal value for γ of about 30°. This gives a very close approximation to the results of the SHT.

Evaluating the performance of any feature extraction routine on real data taken from a complex environment is difficult. The idea of 'correct results' for a line detector cannot be defined without some reference to an intended application. This can be clearly seen when considering how to interpret a curve. Depending on the intended application, one may wish either to ignore it, or to approximate it with straight lines.
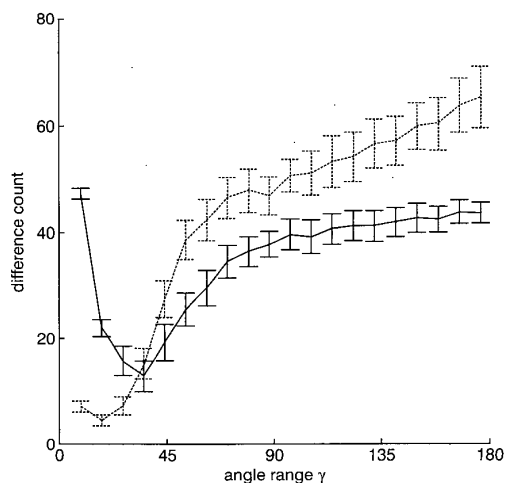
For this reason we discard the abstract concept of correctness and replace it with repeatability. Repeatability is of direct relevance in tracking and object recognition. It is necessary (but not sufficient) that the line extraction routine gives a consistent interpretation of its input data. A measure of repeatability gives a limit to the best performance you can expect of algorithm that depends on a particular interpretation.

The experiments were carried out on image sequences captured from a CCIR601 source, using the intensity value only. The resolution was halved in both directions by summing pairs of adjacent pixels, from the corresponding positions in the fields of each frame.

Input images were first processed by a Deriche [15] edge detector. Subsequently linear non-maximum suppression with four-connectivity was used to find the edgels, and gradient direction calculated from the Deriche results. Typical edge images from the sequence, from photos in Figs. 7 and 8, can be seen in Figs. 9 and 10. The difference between consecutive images in a sequence is relatively small, but large enough to generate different line interpretations even in a deterministic algorithm like the SHT.

Lines extracted from successive frame were compared. Every pair of lines which have both end points within 5 pixels were considered as matches. Only lines 10 edgels and longer were considered in the comparison. Table 1 gives a typical set of results from such a comparison. $\gamma = 40$ in all experiments.

One of the most striking features of the results shown in Table 1 is the improvement in performance made when edgels with full 360 gradient direction were used. The PPHT on both sets of data showed a 50% improvement in the fraction of lines that were successfully matched. The improvement to the PPHT was greater than that seen in the SHT. This brings the performance of the PPHT very near to that of the SHT, within 10%. Table 2 shows the execution time. The implementation for the SHT was kept identical, where possible, with that used for the PPHT. The one area of the SHT that was not optimised was the search of a peak in the accumulator space. Even allowing that it could be speeded up many times, the PPHT far outperforms the SHT.
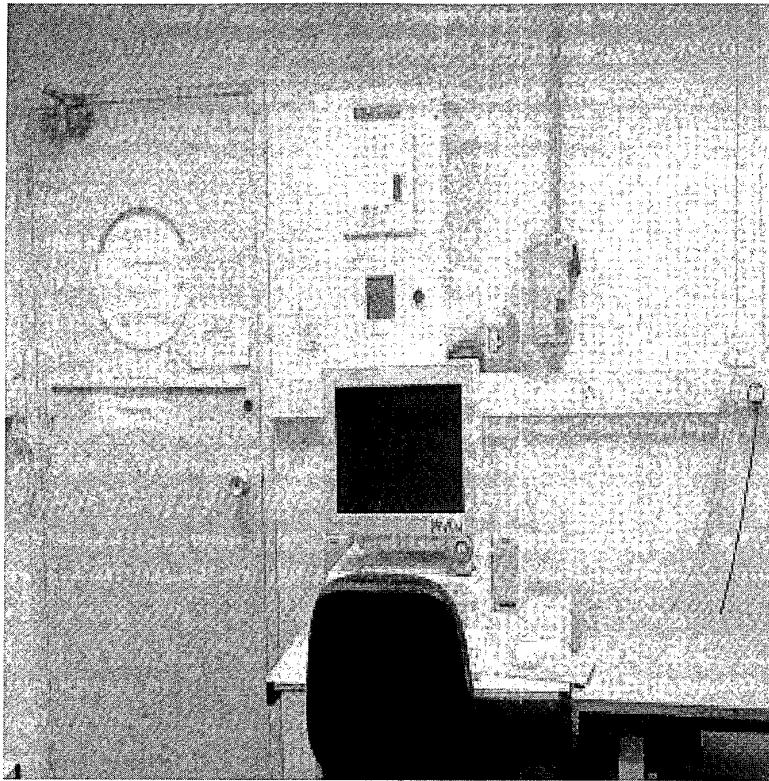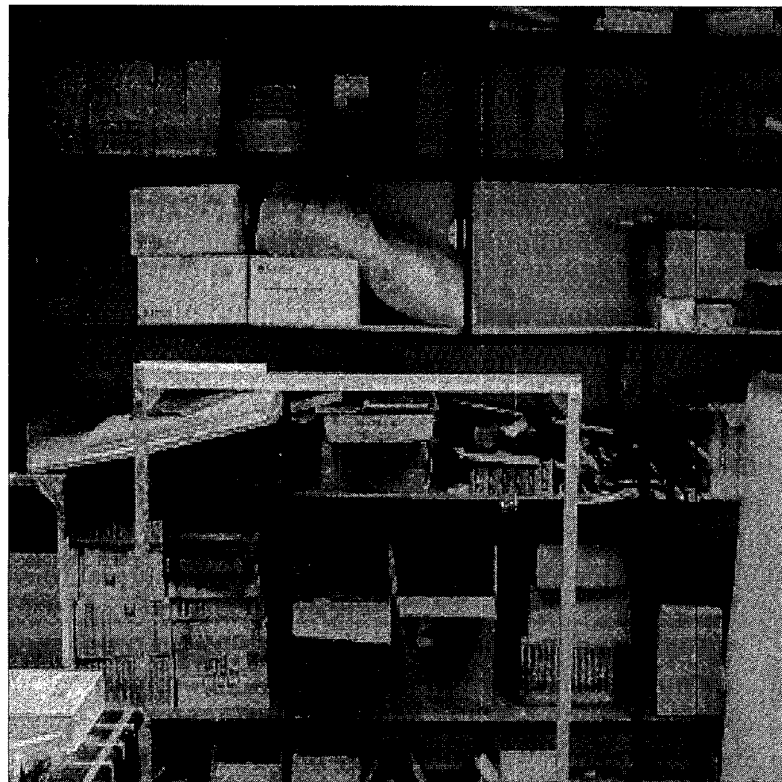
Fig. 7 *Typical door image*
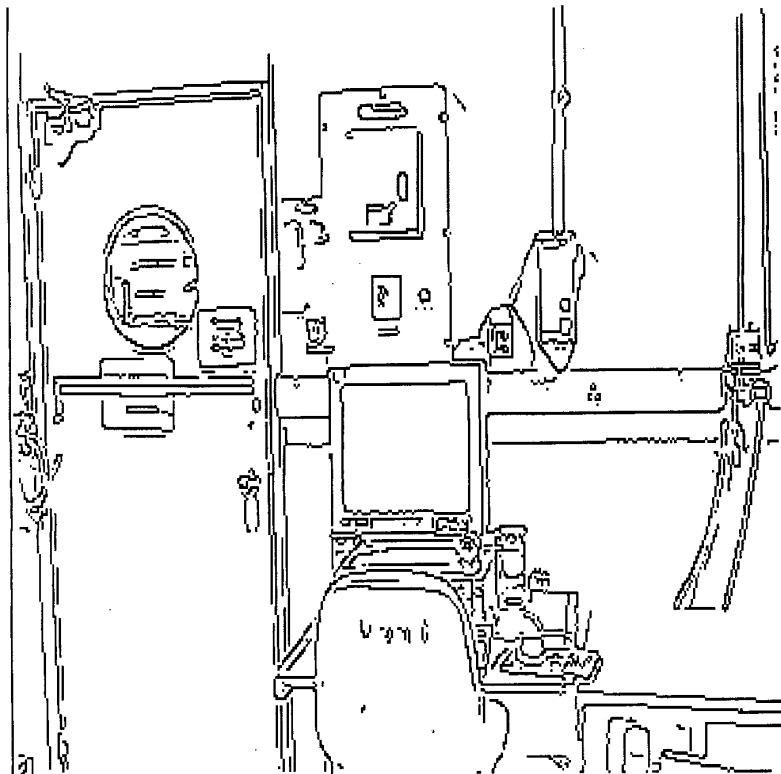

Fig. 8 *Typical shelves image*
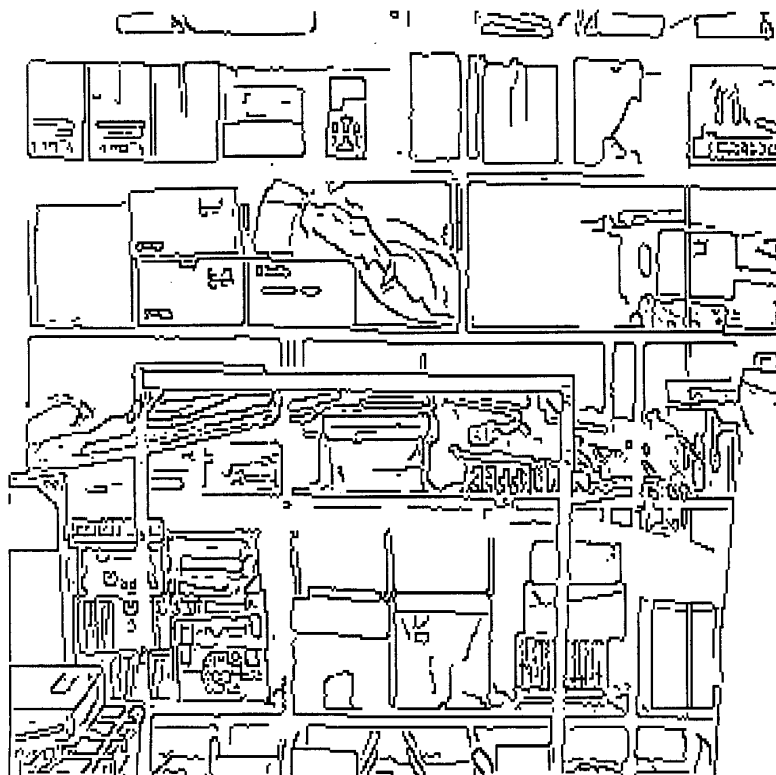
Fig. 9 *Example door edge images*

Fig. 10 *Example shelf edge images*
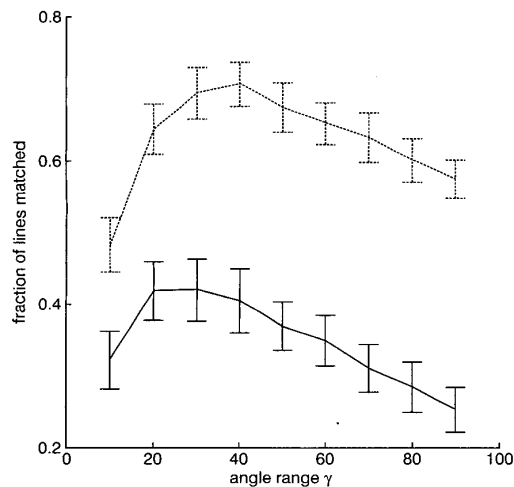
Table 1: Comparison of the PPHT and the SHT

Data set	Method	Range	Matches	Std. Dev.	Ratio	Std. Dev.	Lines	Std. Dev.
door	PPHT	180	57.42	6.32	0.41	0.044	144	5.24
door	SHT	180	110.06	5.63	0.60	0.031	187	4.04
door	PPHT	360	111.89	4.70	0.71	0.030	160	3.16
door	SHT	360	145.79	4.45	0.77	0.026	192	3.20
shelves	PPHT	180	130.51	7.26	0.43	0.025	310	5.82
shelves	SHT	180	207.50	7.66	0.61	0.024	341	5.51
shelves	PPHT	360	211.60	7.92	0.67	0.024	320	4.82
shelves	SHT	360	257.11	6.81	0.76	0.018	342	4.82

Table 2: Execution times for processing 100 frames

Data	Method	Range	Time
door	PPHT	180	55
door	SHT	180	590
door	PPHT	360	60
door	SHT	360	1142
shelves	PPHT	180	148
shelves	SHT	180	1349
shelves	PPHT	360	150
shelves	SHT	360	2297

It is interesting to note that there is a little extra computational cost in using full 360° direction information (the differences arise because of the effects a larger accumulator on memory caching). In the SHT, the increased accumulator size doubles the cost of scanning for peaks, while in the PPHT, because the scan is done during voting, the extra cost is negligible.

Figs. 11 and 12 show the effect of using gradient direction on the stability of the sequence of real images. These results agree well with those given on the synthetic data. The best performance for the door image was with γ set at 40° and for the shelves the value was around 30. These figures correspond well with the minimum in

Fig. 12 *Fraction of lines matched as a function of angle width for the shelves data*

—— 180° edgels
– – – 360° edgels

false negatives and positives seen in Fig. 6. They also support the results on the synthetic data, shown in Figs. 4 and 5.

5 Discussion

The results of the experiments clearly show an improvement in performance for the PPHT. They show that when the PPHT is used with gradient direction information, it has a performance similar to the SHT, even where the SHT uses the same information.

From Figs. 4 and 5, it can also be seen that algorithm is robust with respect to the confidence angle interval. When the range is set too low, or the uncertainty increases the PPHT uses more votes to compensate for this missing information. This is important if the uncertainty in the angles varies. It allows this parameter to be set at an optimal value, without fear that the algorithm will completely fail if angles become more noisy than usual for a short while.

The results on real images summarised in Figs. 11 and 12 demonstrate that the benefits shown on synthetic data extend to the processing of real image data. With the added stability in the generated output achieved by the use of gradient direction, the PPHT becomes an attractive choice for real-time processing of edge images.

Fig. 11 *Fraction of lines matched as a function of angle width for the door data*

—— 180° edgels
– – – 360° edgels

6 Conclusion

The simple modifications shown here notably improve both the accuracy and performance of the PPHT, when gradient direction is available. It has also been shown that the relative improvement in the PPHT is significantly greater than that seen with equivalent modifications to the SHT.

Even where gradient direction is not directly available, it is possible to use neighbouring edgels to estimate the required information successfully. The main disadvantage of using gradient direction is the addition of an extra parameter that defines the uncertainty in the edgel angles. This, however, can be estimated easily by either tuning to optimise performance or by comparing the angles of the edgels to those of the lines they are finally assigned to. Otherwise, the proposed modifications are easy to implement, and the improvements are gained without any significant drawbacks.

7 References

1 MATAS, J., GALAMBOS, C., and KITTLER, J.: 'Robust detection of lines using progressive probablistic Hough transform', *Comput. Vis. Image Underst.*, 2000, **78**, (1), pp. 119–137
2 GALAMBOS, C., MATAS, J., and KITTLER, J.: 'Progressive probabilistic Hough transform for line detection'. Proceedings of IEEE Computer Society Conference on *Computer vision and pattern recognition*, Los Alamitos, California, June 1999, pp. 554–560
3 MATAS, J., GALAMBOS, C., and KITTLER, J.: 'Progressive probabilistic Hough transform'. Proceedings of British Machine Vision Conference BMVC98, 1998, pp. 256–265
4 KIRYATI, N., ELDAR, Y., and BRUCKSTEIN, A.M.: 'A probabilistic Hough transform', *Pattern Recognit.*, 1991, **24**, (4), pp. 303–316
5 YLA-JAASKI, A., and KIRYATI, N.: 'Adaptive termination of voting in the probabilistic circular Hough transform', *IEEE Trans. Pattern Anal. Mach. Intell.*, 1994, **16**, (9), pp. 911–915
6 YLA-JAASKI, A., and KIRYATI, N.: 'Automatic termination rules for probabilistic Hough algorithms'. Proceedings of 8th Scandinavian Conference on *Image Analysis*, 1993, pp. 121–128
7 SHAKED, D., YARON, O., and KIRYATI, N.: 'Deriving stopping rules for the probabilistic Hough transform by sequential-analysis', *Comput. Vis. Image Underst.*, 1996, **63**, (3), pp. 512–526
8 O'GORMAN, F., and CLOWES, M.B.: 'Finding picture edges through collinearity of feature points', *IEEE Trans. Comput.*, 1976, **C-25**, (4), pp. 449–456
9 PALMER, P.L., KITTLER, J., and PETROU, N.: 'Using focus of attention with the Hough transform for accurate line parameter estimation', *Pattern Recognit.*, 1994, **27**, pp. 1127–1133
10 ILLINGWORTH, J., and KITTLER, J.: 'A survey of the Hough transform', *Comput. Vis. Graph. Image Process.*, 1988, **44**, pp. 87–116
11 LEAVERS, V.F.: 'Which Hough transform', *CVGIP, Image Underst.*, 1993, **58**, (2), pp. 250–264
12 KALVIAINEN, H., HIRVONEN, P., XU, L., and OJA, E.: 'Probabilistic and nonprobabilistic Hough transforms—overview and comparisons', *Image Vis. Comput.*, 1995, **13**, (4), pp. 239–252
13 DUDA, R.E., and HART, P.E.: 'Pattern classification and scene analysis' (John Wiley, 1973)
14 PRESS, W.H., FLANNERY, B.P., TEUKOLSKY, S.A., and VETTERLING, W.T.: 'Numerical recipes in C' (Cambridge University Press, 1992)
15 DERICHE, R.: 'Using Canny's criteria to derive a recursively implemented optimal edge detector', *Int. J. Comput. Vis.*, 1987, **1**, pp. 167–187