

Using gradient information to enhance the Progressive Probabilistic Hough Transform

C. Galambos, J. Kittler

CVSSP,

University of Surrey,

Guildford, Surrey GU2 5XH,

United Kingdom

Email: C.Galambos@ee.surrey.ac.uk

J. Matas

'CVSSP' and 'Centre for Machine Perception',

Czech Technical University,

Karlovo náměstí 13, 12135 Praha,

Czech Republic

Abstract

In this paper we look at the benefits to be gained in using gradient information to enhance the Progressive Probabilistic Hough Transform (PPHT). It is shown how using the angle information in controlling the voting process and in assigning pixels correctly to a line, PPHT's performance can be significantly improved. The improved algorithm gives results very close to that of the Standard Hough Transform, but requires significantly less computation.

1 Introduction

In this paper we investigate the benefits that can be gained by using gradient information in conjunction with the *Progressive Probabilistic Hough Transform* (PPHT) algorithm [5] [1].

PPHT is a variant of the Standard Hough Transform (SHT) algorithm for line detection. It is based on the Probabilistic Hough Transform (PHT) [4]. (See [3] for an overview of these algorithms). PPHT differs from PHT in that the accumulator space is constantly scanned for significant peaks and lines are extracted as they are found. When a line is detected, all pixels that are assigned to the line are either removed from the list of unused pixels or if they have voted their votes are removed from the accumulator space.

PPHT has several advantages; it removes the need for apriori knowledge about the lines lengths in the image; it is well suited to recovering many lines from a single image and it can be stopped at any time whilst still giving useful results. These properties make it ideal for real-time applications like tracking. Details and justification for these claims can be found in [1].

The only parameter PPHT requires to be set is the *false positive threshold* (l). This controls the number of acceptable false positives that can be generated by the algorithm. Typical values range from 0.1 to 10^{-9} . The effect of this

parameter on the total number of votes required to process an image is presented in [1]. In essence, the lower 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 consider the use of extra information to improve the performance of PPHT. Though similar improvements are possible with SHT they have been considered elsewhere [2] therefore their performance benefits will not be examined in detail here.

2 Modifications to the PPHT algorithm

In using gradient information we make two key changes to the standard PPHT algorithm. First we limit the range of angles voted for by each pixel. Secondly when a peak is detected and the pixels are being assigned to the line we check whether the gradient of each pixel is consistent with the angle indicated by the peak.

Limiting the range of angles has two main effects. The first is to reduce the computation required to process a new pixel. The second is to reduce the clutter in the Hough space, or in other words, increase the signal to noise ratio. Gradient information can be made available from several sources. Ordinary edge detectors give this information for free. If edge information is not available directly, it is possible to estimate it by calculating the moments of the neighbouring edge pixels.

It has been shown that the cost of voting into the Hough space is a major factor in the computation cost of the algorithm [5]. Cutting the number of angles voted for proportionally reduces the cost of casting the vote. As long as the angles considered in voting reasonably cover the uncertainty in the angle of the gradient, then no information is lost, the gain in computation speed is free.

The second advantage of restricting the angle is the reduced 'clutter' in the Hough space. This significantly reduces the chances of a false positive peak occurring. If one

considers that only one angle is correct out of the many bins voted in, halving the angles reduces the 'noise' votes by nearly half. Again the limit on how far the angle may be reduced is dictated by the uncertainty in the gradient of the edgel.

In addition the gradient information can be used again in the post processing part of PPHT. In normal PPHT after a peak is detected a search is made in the image space for all pixels that could correspond to the line. This search can be refined with PPHT to only include pixels with matching gradients. The same range of angles is used for this as for voting. This can significantly improve the quality of lines generated, both in pixel assignment at 'T' junctions, and where there are many closely spaced lines.

The key factor in determining how much improvement is gained is the uncertainty of the gradient of each edgel. Unfortunately this is very difficult to quantify theoretically, and even if it were possible, the results would only be applicable to a particular edgel source. For the purposes of this paper we will assume the angular error distribution is unknown but approximately uniform for each source.

It is also conceivable that one may be able to estimate the uncertainty for each edgel, and incorporate that into the algorithm described here, but its not clear that this would be possible for all sources, and so we will only consider the simpler case here.

Finally there are two possible ranges of angles that can be considered; either 180 or 360 degrees. In the SHT, normally only a 180 degree range is considered in the accumulator space, as there is no way to distinguish between lines running in opposite directions. However, in cases where a full 360 degree angle is available it is possible to double the size of the accumulator space, whilst only voting for the same number of points. This halves the density of 'noise' votes and allows the use of lower thresholds. The added angular information also resolves some ambiguities that can arise from close parallel lines thus improving results further.

Another possible refinement of the voting process would be to move from a simple integer scheme, to one where the value added to each accumulator space represented the probability of it being the correct edgel given the measured angle and the uncertainty [6]. To do this properly however, one would also have to vote proportionally in adjacent ρ bins, thus significantly increasing the computational cost of voting. Though this may well lead to improved performance, the benefits are not as clear as the effects of the proposed improvements and so this idea will not be explored here.

3 Experiments

For consistency the stopping criterion of the algorithm exploiting gradient information should be modified to take




Figure 1. Example of a synthetic edge image.

account of the reduced noise in the accumulator statistics. In order to make clear performance comparisons between the modified and unmodified PPHT, this has not been done. There are several approximations made when calculating the threshold. Changing them at the same time would make it difficult to pin down the magnitude of the performance improvements. In the following experiments, results with an angle of 180 degrees are equivalent to the original methods.

The synthetic images used for these experiments were 256 pixels squared, each with 20 lines of random length between 1 and 100 pixels. An example of such a synthetic image can be seen in Figure 1. Each experiment was repeated 100 times. The means and the standard deviations are shown in the graphs. To make the gradient information in the synthetic image realistic it was estimated by counting the moments of all the edge pixels within a radius of 2.5 grid squares. This estimation works well on real images as well and sometimes better as they often contain fewer crossing lines.

Figure 2 shows the number of voting operations used in processing the images. Though voting for a restricted range 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 pixels, and hence the full SHT uses 2000 voting operations. This means the results for PPHT were obtained with about one tenth of the operations required by SHT.

As the range of angles used drops below the uncertainty in the gradient information the number of votes needed to process the image starts to rise as can be seen in Figure 2 for angles below 30 degrees. This occurs because the votes cast in the accumulator space sometimes fail to cover the actual line position, reducing the size of the peaks and forcing the PPHT to cast more votes to detect them.

In Figures 4 and 3 we show the detection performance results for PPHT as a function of the gradient angle constraint. The performance is measured in terms of average numbers of false positives (Figure 3) and false negatives (Figure 4)




Figure 2. PPHT Votes, using gradient info for accumulation.




Figure 3. PPHT False positives, using gradient info for accumulation.

(undetected lines) as compared with the know ground truth for each test image. We note that by restricting the range of edgel orientation angles to an interval of 30 - 60 degrees the number of false negatives dips significantly for the faster version of 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 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 reasonable flat between 30 and




Figure 4. PPHT False Negatives, using gradient info for accumulation.

60 degrees 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.

When compared with the results for SHT as seen in Figure 5, for $l=10^{-1}$ the number of false negatives is about the same as for PPHT for angle ranges between 30 to 70 degrees. There are about twice as many false positives for PPHT over SHT. Considering the relative execution times this is a reasonable result.

The results are shown in Figures 3 and 4. With l set to 10^{-4} the results for false positives are fairly similar for both PPHT and SHT, but the number of false negatives shows about a 20 percent drop over the previous results for angles 30 to 60 degrees. This means the overall results for PPHT, at least on synthetic images are better than those for SHT. As can be seen from Figure 2 this extension has little impact on the number of voting operations required to process the image. Since SHT does not benefit as much as PPHT from the use of gradient information, we can conclude the added information helps to make up for that which is lost in the process of sub sampling the input edgels.

In the final experiment we compare the output of SHT and PPHT on a real image. There is insufficient space in this paper for full comparison, so this experiment is only intended to illustrate that the improvements give similar benefits when processing real image data. These experiments were run with the house edge image as used in [5]. It is worth noting that PPHT with gradient information has been tested successfully with many other real images.

The SHT was run with angle range of 30 degrees to give as near optimal interpretation as possible. SHT and PPHT use different stopping rules and hence the number of short




Figure 5. SHT with gradient information.

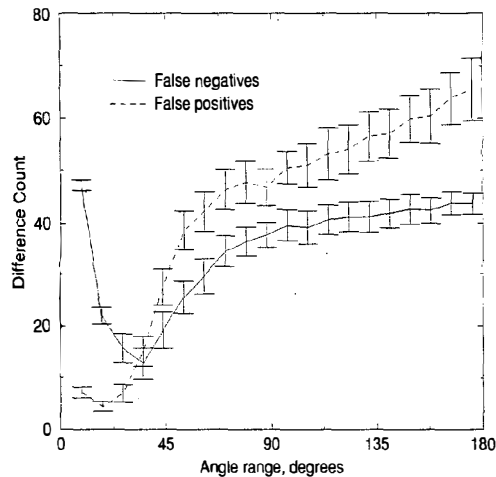


Figure 6. Comparison of PPHT and SHT with gradient information.

lines recovered vary significantly. To reduce problems with this causing excessive false positives only lines of 10 pixels and longer were used in the comparison. PPHT was run with an l of 10^{-4} which has been found to give good performance.

Figure 6 shows the results of the comparison. As indicated by the experiments on synthetic data there is an optimal angle of about 30 degree's. This gives a very close approximation to the the results of SHT.

4 Discussion

The results of the the experiment clearly show an improvement in performance for PPHT. They show that when

PPHT is used with gradient information it has a performance similar to SHT, even where SHT uses the same information. Since the gradient information was estimated using neighbouring pixels, comparable error bounds on gradient information estimates can always be achieved for any image data.

From Figures 3 and 4 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 PPHT uses more votes to compensate for this missing information. This is important if the uncertainty in the angles varies. It allows the value of 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.

5 Conclusion

The simple improvements show here significantly improve both the accuracy and performance of PPHT, where gradient information is available. Even where it is not, using moment's to estimate the information can lead to gains in accuracy. The main disadvantage is the addition of an extra parameter which gives 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.

References

- [1] C. Galambos, J. Matas, and J. Kittler. Progressive probabilistic Hough Transform for line detection. In *Computer Vision and Pattern Recognition*, pages 554–560, Los Alamitos, California, June 1999. IEEE Computer Society.
- [2] J. Illingworth and J. Kittler. A survey of the Hough Transform. *Computer Vision, Graphics and Image Processing*, 44:87–116, 1988.
- [3] H. Kalviainen, P. Hirvonen, L. Xu, and E. Oja. Probabilistic and nonprobabilistic Hough Transforms - overview and comparisons. *Image And Vision Computing*, 13(4):239–252, 1995.
- [4] N. Kiryati, Y. Eldar, and A. M. Bruckstein. A probabilistic Hough Transform. *Pattern Recognition*, 24(4):303–316, 1991.
- [5] J. Matas, C. Galambos, and J. Kittler. Progressive probabilistic Hough Transform. In M. S. Nixon, editor, *Proc British Machine Vision Conference BMVC98*, volume 1, pages 256–265, London, UK, September 1998. University of Southampton, British Machine Vision Association.
- [6] P. L. Palmer, J. Kittler, and M. Petrou. Using focus of attention with the Hough Transform for accurate line parameter estimation. *Pattern Recognition*, 27:1127–1133, 9 1994.