Edge Detection at Junctions

Du Li, G.D. Sullivan & K.D. Baker

Intelligent Systems Group, Department of Computer Science,
University of Reading, Reading RG6 2AX, UK
Du.Li@reading.ac.uk

This paper discusses the poor performance of the Canny edge detector at junctions formed by three edges, where connectivity is usually destroyed in an arbitrary way. This problem is shown to be due to the non-maximal suppression algorithm, used to identify points of inflection in the grey-level surface. A method to overcome the problem, based on local evidence in the gradient vector, is reported and demonstrated.

1. INTRODUCTION

Edge detection has attracted much attention in the field of computer vision. It is a commonly used step towards a symbolic description of the image. A number of edge detectors have been proposed, including that by Canny [1,2], which has received widespread practical acceptance.

Edges are usually defined as being at the points of inflection of the grey level surface, and the inflection points may be identified by finding the local maxima in the gradient of the image. In the usual implementation of the Canny edge detector, the maxima are found by means of an algorithm, which suppresses all points where the magnitude of the gradient is not locally maximal in the direction of the gradient at the point in question. We refer to this algorithm as the conventional Canny operator.

However, junctions of 3 or more linear edges are not treated effectively. Such junctions are very common in images, especially at trihedral intersections and occlusion boundaries. The failure of the Canny operator at junctions destroys the connectivity between edges in the image, which may cause problems in later grouping stages. Figure 1 shows two synthetic T-junctions formed by the edges between three homogenous grey level patches. The conventional Canny operator leaves a gap on the weaker edge (Figure 2). Depending on the relative strengths of the edges, the Canny operator breaks a single real edge into two pieces (Figure 2(a)) or fails to detect the junction with the occluded line (Figure 2(b)).

T-junctions are powerful indicators of 3-d occlusion [3, 4, 5] which the Canny operator fails to identify. Similar problems occur at trihedral (Y or ➔) junctions. These failures may not seem significant to intelligent human eyes, but they create problems for higher level algorithms concerned with the connection and grouping of edges. We show below that the problem is caused by the non-maximum suppression algorithm, which works effectively at linear edge features but fails at junctions. We call such failures "false suppression".

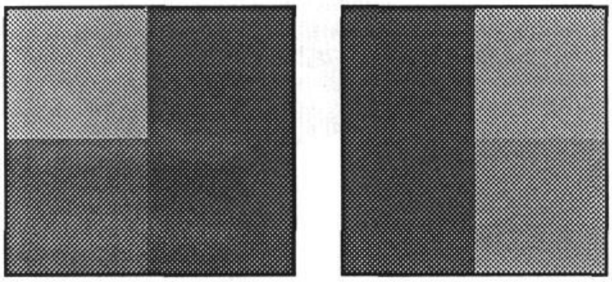

Figure 1. T-Junctions formed by 3 regions of different grey levels

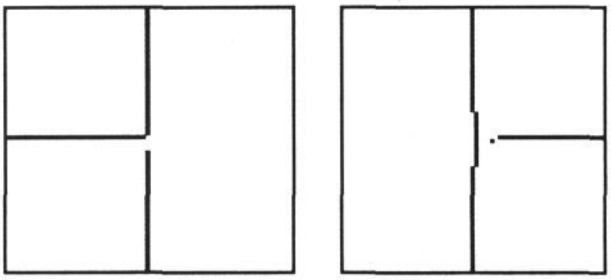

Figure 2. The false suppression of edges by the Canny operator

The Canny operator has previously attracted much critical evaluation, e.g. [6, 7, 8, 9, 10]. A few studies have considered the failure of connectivity at junctions in detail. Nobel [9] criticised this behaviour of the Canny operator at junctions and proposed a completely different approach, based on mathematical morphology to overcome the problem. It is more common to assume that such failures can be corrected at a later stage of analysis, by use of higher level knowledge [10]. The latter approach inevitably makes the search at the line grouping stages of analysis more ambiguous. This uncertainty can be avoided if the junctions are detected properly.

This paper examines the reason for failure at junctions and suggests a locally based algorithm to overcome the shortcoming. Finally we discuss the limitation of the proposed solution and put forward proposals for further work.

AVC 1989 doi:10.5244/C.3.21

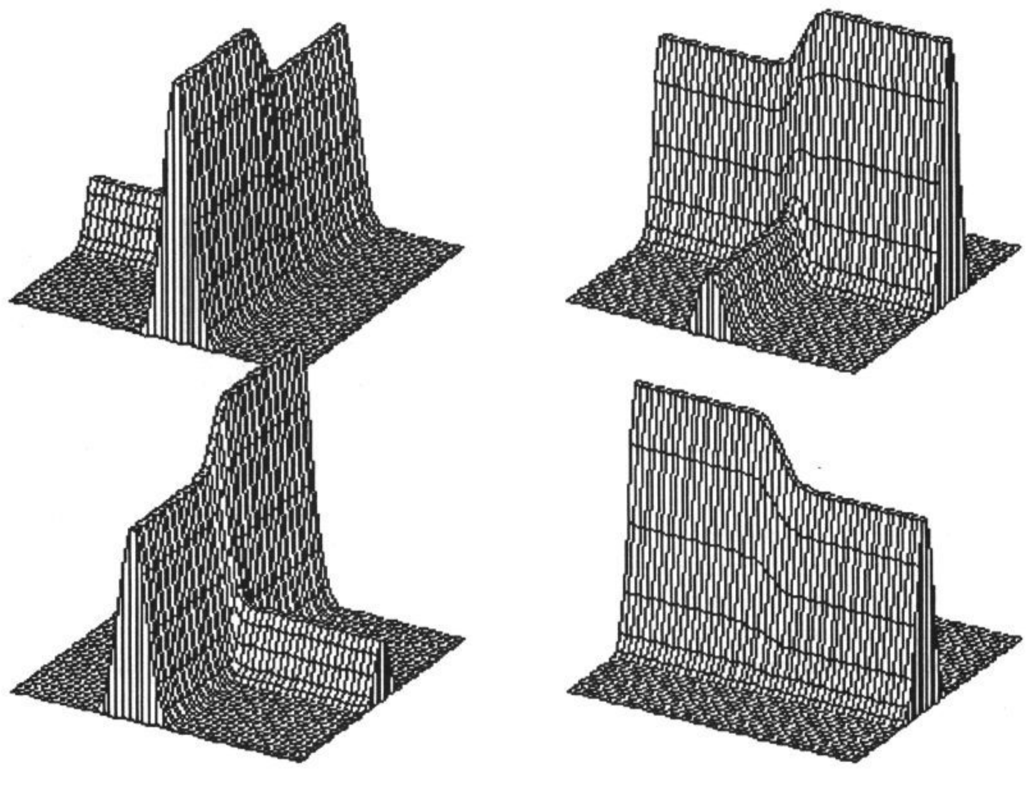

Figure 3. Four views of the magnitude of ∇I_g. The weaker ridge coincides with the positive X-axis

THE PERFORMANCE OF THE CANNY OPERATOR AT JUNCTIONS

We consider a T-junction formed by 3 homogenous grey level surface patches (Figure 1(b)). The intensity surface is represented by $I(x,y)$.

Edges are identified by the Canny operator as points of inflection in I, after smoothing by a gaussian of standard deviation σ, to form I_g ($\sigma = 2$ in the Figures). Figure 3 illustrates the magnitude of the discretely sampled gradient vector, ∇I_g. This consists of distinct ridges along the edges in the image, which merge at the junction.

Figure 4 represents the gradient vector in the immediate vicinity of the junction by means of directed lines. It can be seen that the gradient direction lies perpendicular to the edges at points which lies further than approximately 2-3σ from the junction, but becomes distorted by interference between the edges close to the junction.

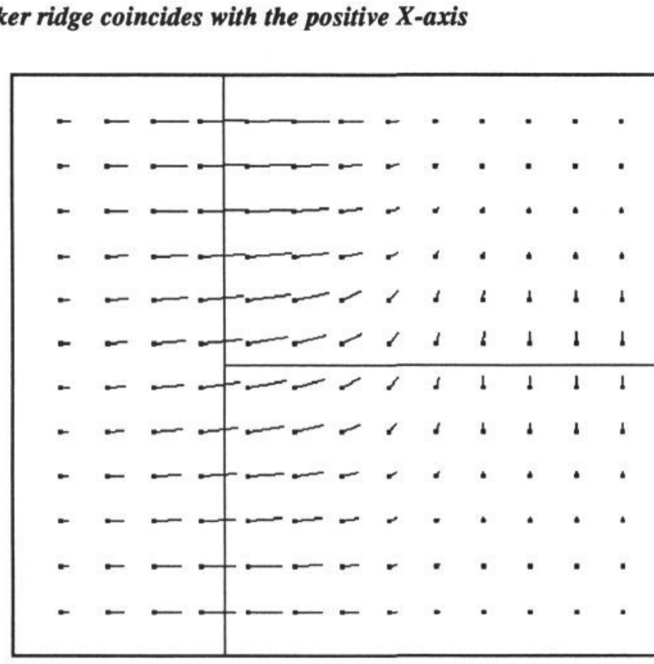

Figure 4. Gradient vectors in the neighborhood of a T-junction

The Canny operator discovers 1-dimensional maxima in the magnitude of the gradient vector, by computing the change in the magnitude of the gradient at each pixel, in the direction of the gradient. It does this by interpolating values between samples, at a distances of approximately +/-1 pixel from each point. As we approach the junction along the weakest edge (the horizontal edge in Figure 4), the local direction of ∇I_g

becomes increasingly affected by the vertical edge. At some point, depending on the relative grey levels of the image, it becomes rotated towards the x direction. The magnitude of ∇I_g is no longer maximal (as defined by the Canny operator), since it increases in the x direction towards the vertical ridge. The non-maximum suppression algorithm therefore rejects this part of the weaker edge, and leaves a gap on the weaker horizontal edge, as demonstrated in Figure 2(b).

RECOVERY FROM FALSE SUPPRESSION

Identification of false suppression

According to the preceding analysis, false suppression reliably occurs on the weakest edge of a junction, where a false free end is created. "Real" free ends do occur in images, where a grey edge tapers out. We can distinguish real from false free ends by the fact that at a false free end the magnitude of the gradient increases in the general direction perpendicular to the gradient at the free end, and a new peak occurs within approximately 3σ.

Recovery Algorithms

A simple algorithm to correct the error is to extend the falsely suppressed edge perpendicular to the gradient at the free end, until it reaches a point which has already been identified as an edge by the non-maximum suppression algorithm. To be acceptable, the magnitude of the gradient at the new edge must exceed that of the free end.

In practice the direction of the gradient at the free end is unreliable, and our algorithm searches back along the affected edge for a distance of 2σ to check that the edge is not an isolated fragment, and that it has a consistent orientation and strength.

As shown in Figure 5 situations may occur where simple linear extrapolation causes an error in the recovery of the junction position. However, the correctly detected stronger edges of the occluding surface are necessarily of different strengths. Therefore an inflection point should exist in the magnitude of the gradient along the already marked stronger edges. This may be seen clearly in Figure 3.

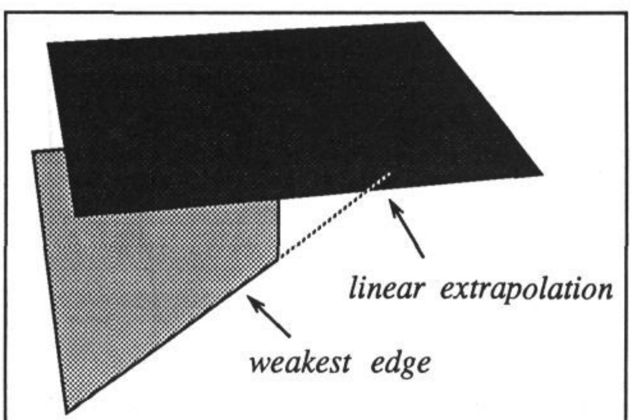

Figure 5. Failure of the linear algorithm

A more accurate algorithm searches for the inflection point on the strong edge, and connects the free end to this point. Figure 6 shows the algorithm applied to Fig. 1. In fact the location of the point of inflection still fails to identify the junction point exactly, since it may be displaced by the gaussian smoothing operation, up to a distance of σ from the expected position. Information to correct this residual error is not available in the blurred image at this scale of processing.

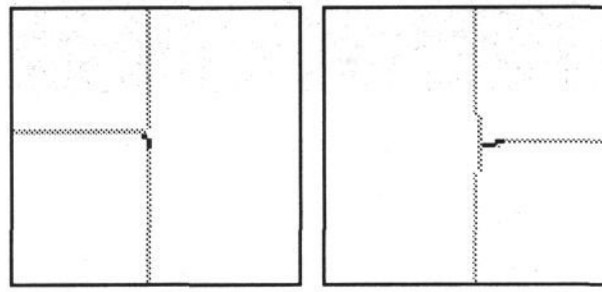

Figure 6. Recovery of edge-connectivity by the modified operator

Application to natural images

The previous discussion is based on the behaviour at ideal edges, as was Canny's original treatment. In natural images edges make junctions with arbitrary orientations, the regions are not homogeneous, and the image is degraded by noise.

Figure 7 illustrates the performance of the algorithm applied to a natural image. Figure 7(a) shows the image, Figure 7(b) the result of the conventional Canny operator and Figure 7(c) the result of modified operator. Attention should be focussed on the junctions. In Figure 7(c) the grey points are those marked by the conventional Canny operator and black points are the additional points marked by the current method. Note the improved ability to detect occluded edges, especially at the windows of the rear car. These T-junctions have now been made explicit, and this helps in the interpretation of the occlusion relationship between the cars.

Discussion.

We have demonstrated a modification to the conventional Canny operator, which overcomes the failure to detect junctions, by a purely local analysis of the gradient vector.

It has previously been suggested that recovery of connectivity should be left to higher stages of vision, whilst deriving the symbolic description of the edges, or during subsequent grouping stages. There are two arguments against this approach.

Firstly, it is unwise in principle to resort to higher level constraints when information is available at a lower level. For example a connected sequence of edges may contain a dog-leg which is too insignificant to be identified on its own, but which may be correctly segmented if a T-junction is detected at one of the corners. An example occurs in Figure 7 where the rear

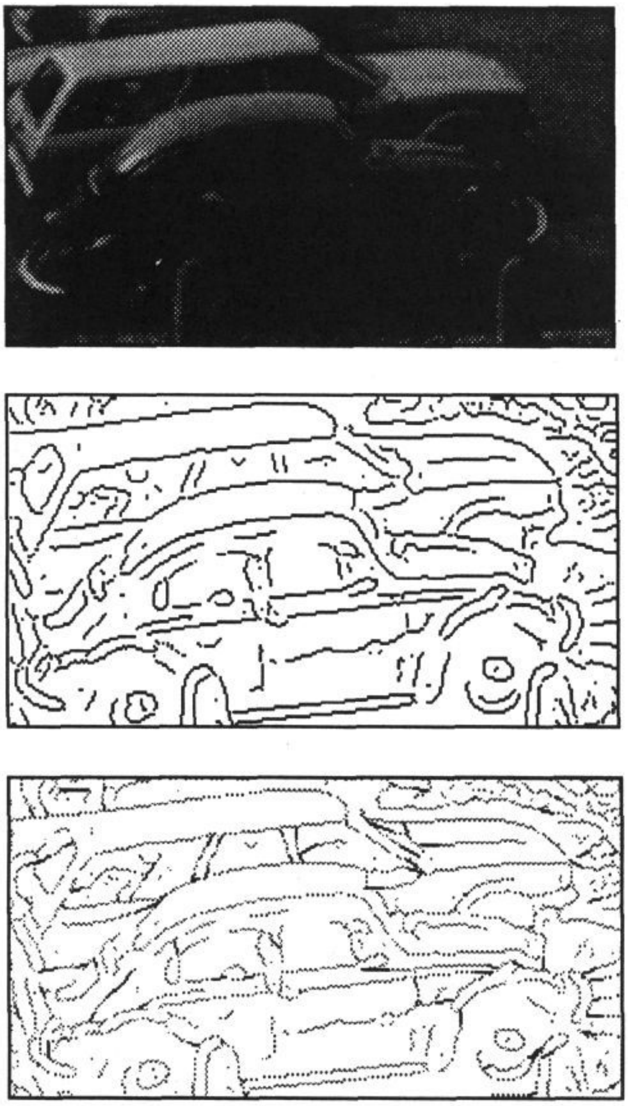

Figure 7. ***(a) Original Image***
(b) Results of the Canny operator
(c) Results of the modified operator

roof of the VW runs into the window of the Estate car Figure 7(b), but is connected along the back of the VW in Figure 7(c).

Secondly, any additional ambiguity in the line connectivity delivered by low-level processing necessarilly increases the size of the search problem encountered in later grouping stages.

Figure 8(a) illustrates the latter point by means of a synthetic "Mondrian" image. The Canny operator (Figure 8(b)) picks out connected sequences of edgels which wander haphazardly along unrelated edges between blocks of colour. It is difficult to segment the blocks of colour on the basis of the edge evidence alone. The modified operator (Figure 8(c)) successfully locates most of the T-junctions in the image, and only fails at junctions involving very short edges. This greatly facilitates the discovery of closed polygons, which are the boundaries of the separate blocks.

Our demonstrations show that there is information in the smoothed gradient vector which is not used by the Canny operator. Edge detectors are commonly used in

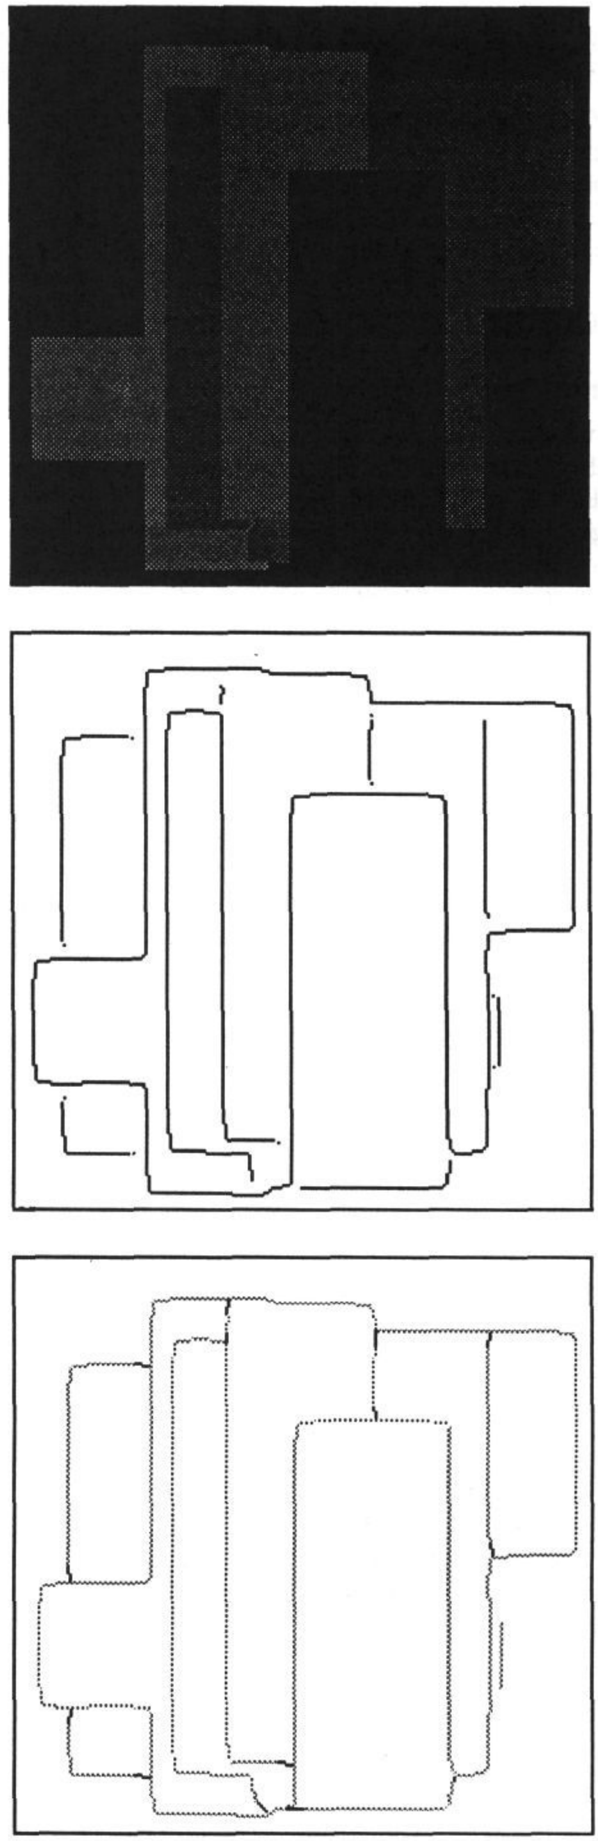

Figure 8. ***(a) Original Image***
(b) Results of the Canny operator
(c) Results of the modified operator

model-based vision to identify focus features [11] or cues [12], which invoke hypotheses of known objects. In our experience, this stage of feature analysis is bedevilled by the imprecision of the low-level processing. Any improvement in the data-driven extraction of features is greatly to be desired.

References.

1. Canny, J. A Computational Approach to Edge Detection, IEEE PAMI Vol. 8, No. 1, pp679-698, 1986

2. Canny, J. Finding Edges and Lines in Images, Report AI-TR-720, Artificial Intelligence Laboratory, Massachusetts, 1983

3. Guzman, A. Decomposition of a Visual Field into Three Dimensional Bodies, *Automatic Interpretation and Classification of Images* (ed. A. Grasselli), pp243-276, 1969

4. Huffman, D. A. Impossible Objects as Nonsense sentences, *Machine Intelligence 6* (ed. Meltz, B. and Michie, D.), 1971

5. Clowes, M. B. On Seeing Things, Artificial Intelligence, Vol. 2, No. 1, pp79-116, 1971

6. Spacek, L. Edge Detection and Motion Detection, Image and Vision Computing, Vol. 4, No. 1, 1986

7. Petrou, M. On the Optimal Edge Detector, Proceedings of AVC88, pp191-196, 1988

8. Nobel, J. A. Finding Corners, Proceedings of AVC87, pp267-274, 1987

9. Nobel, J. A. Morphological feature Detection, Proceedings of AVC88, pp203-210, 1988

10. Rake, S. Finding Curved Ribbons, IBM report UKSC 200, 1988

11. Bolles, R. C. Recognizing and Locating Partially Visible Objects: The Local-Feature-Focus Method, The International Journal of Robotics Research, Vol. 1, No. 3, 1982

12. Bodington, R. Baker, K. D. and Sullivan, G. D. The Consistent Labelling of Image Features Using an ATMS, Proceedings of AVC88, pp7-12, 1988