EXTRACTING STRAIGHT LINES

J. Brian Burns, Allen R. Hanson, Edward M. Riseman

COINS TECHNICAL REPORT 84-29

Reprinted from IEEE SEVENTH INTERNATIONAL CONFERENCE ON PATTERN RECOGNITION Proceedings
July 30-August 2, 1984 Montreal, Canada

EXTRACTING STRAIGHT LINES

J. Brian Burns, Allen R. Hanson, Edward M. Riseman

Computer and Information Science Department
University of Massachusetts
Amherst, Massachusetts 01003

ABSTRACT

This paper presents a new approach to the extraction of straight lines in intensity images. Pixels are grouped into edge support regions of similar gradient orientation, and then the structure of the associated intensity surface is used to determine the location and properties of the edge. The resulting regions and extracted parameters form two separate representations of a straight line segment, pixel-based and symbolic, that can be used together for a variety of purposes. The algorithm appears to be far more effective than previous techniques for two key reasons: 1) the gradient orientation is used as the initial organizing criterion in the extraction of straight lines, instead of the gradient magnitude; and 2) data directed organization of the complete context of a straight line is determined prior to any local decisions about participating edge points.

1.0 INTRODUCTION

The organization of significant local intensity changes into the more global abstractions called "lines" or "global intensity boundaries" is an early, but important, step in the transformation of the visual signal into useful intermediate constructs for interpretation processes. Despite the large amount of research appearing in the literature, effective extraction of linear boundaries has remained a difficult problem in many image domains. There are two goals of this paper: a) the development of mechanisms for extracting straight lines from complex images, including intensity discontinuities of arbitrarily low contrast; and b) the construction of an intermediate representation of edge/line information which allows high-level mechanisms efficient access to relevant lines. A more detailed presentation can be found in [BUR84].

We contend that the major failings of line extraction algorithms are twofold: the relegation of information about edge orientation to a secondary role in the processing, and the lack of a comprehensive global view of the underlying image structure prior to making local decisions about edge features. In most edge and line extraction algorithms, the magnitude of the intensity change has been used in some manner as the dominant measure of importance of the local edge. It is our view that edge orientation carries important information about the spatial extent of the straight line.

The technique presented here was motivated by a need for a straight line extraction method which can find straight lines in reasonably complex images, particularly those lines that are long but not necessarily of high contrast. A key characteristic of the approach presented here that distinguishes it from most previous work is the global organization of the supporting edge context prior to any decisions about the relevance of local intensity changes. An estimate of the local gradient orientation at each pixel is the basis of these first organizing processes. Grouping pixels into edge support regions avoids the plethora of magnitude responses from masks at varying sizes and orientations, as well as unnecessary complexity in the subsequent organizing mechanisms. The approach presented here has its roots in the "gradient collection" process of Hanson et al [HAN80], as well as [EHR78].

Our approach allows the extraction of straight lines despite weaknesses in line clarity due to local edge inconsistencies or deficiencies in width and contrast. It directly addresses the problems associated with the size of the edge operators and determines the extent of support to be given to edges and lines directly from the underlying data.

2.0 A REPRESENTATION AND PROCESS FOR EXTRACTING STRAIGHT LINES

2.1 Overview

There are four basic steps in extracting straight lines:

1. Group pixels into edge-support regions based on similarity of gradient orientation. This allows data directed organization of edge contents without commitment to masks of a particular size.

2. Approximate the intensity surface by a weighted planar fit. The fit is weighted by the gradient magnitude associated with the pixels so that intensities in the steepest part of the edge will dominate.

[1] This work has been supported in part by the Air Force Office of Scientific Research under contract F49620-83-C-0099.

3. Extract attributes from the edge-support region and the plane fit. The attributes extracted include the representative line and its length, contrast, width, location, orientation, and straightness.

4. Filter on the attributes to isolate various image events such as long straight lines; high contrast short lines (heavy texture); low contrast short lines (light texture); and lines at particular orientations and postitions.

2.2 Grouping Pixels into Edge Support Regions Via Gradient Orientation

Figure 1 shows two representative images used to illustrate the process. Figure 2(a) is a surface plot of a 32x32 subimage from another house image; results for the full images are shown in subsequent sections. The vector field drawn in figure 2(b) shows the corresponding gradient image where the length of the vector encodes gradient magnitude. The gradient estimates were formed by convolving the image with two-by-two edge masks (figure 2(b) inset). Note that the sign of the gradient is relevant.

An extremely simple process was employed to group the local gradients into regions on the basis of the orientation estimates. The 360 degree range of gradient directions is arbitrarily partitioned into a small set of regular intervals, say eight 45 degree partitions or sixteen 22.5 degree partitions. If our conjectures about edge orientation are correct, then pixels participating in the edge-support context of a straight line will tend to be in the same edge orientation partitions and adjacent pixels that are not part of a straight line will tend to have different orientations. A simple connected components algorithm can be used to form distinct region labels for groups of adjacent pixels with the same orientation label (Figure 2(c)). Note that in Figure 2(c) the great degree of fragmentation into many small regions of very low gradient magnitude could be grouped into a homogeneous region later, rather than interpreting them as edge elements.

To make the fixed partition technique more sensitive to edges of any orientation, the current algorithm uses two overlapping sets of partitions, with one set rotated a half-partition interval. Thus, if a 45 degree partition is used starting at 0 degrees, then a second set of 45 degree partitions starting at 22.5 is also used. The critical problem of this approach is merging the two representations in such a way that a single edge is principly associated with a single gradient region. The following scheme is used to select such regions for each edge: first the lengths are determined for the regions; then, since each pixel is a member of exactly two regions (one in each segmentation), the pixel decides which one provides the longest interpretation; finally each region counts up the number of pixels within its boundaries that voted for it as opposed to regions of the other segmentation. The support a region is given is the number of votes for it over the total number of pixels in the region. The regions selected are those that have a majority, i.e., the support is greater than 50%. For further discussions on grouping see [BUR84].

2.3 Interpreting the Edge-Support Region as a Straight Line

The underlying intensity surface of each gradient region is a candidate for a straight line structure; the key problem is to use this information to find the line. The positional parameters extracted will serve as the core of the structure's symbolic description as well as a coordinate system about which other attributes will be measured.

In this section, we will examine a simple process for computing the parameters of a planar fit to the intensity surface of the pixels in each edge-support region. The region depicted in figure 3(a) and as dots in the surface plot of figure 2(a) will serve as our example. Note that it includes pixels outside the group of gradients depicted in figure 2(c), since the two-by-two masks incorporated them in the gradient estimation. Haralick [HAR81] modelled the local intensity surface in the neighborhood of a pixel as a planar surface patch called a 'sloped facet'. This planar fit served as a model of the region structure and was used to determine if the pixel was at a region boundary or not. In our application the planar fit will be applied to all pixels in a support region instead of an a priori fixed geometric configuration. If a direct least-square planar fit to all pixels in the support region is computed, then many pixels which might be at the tail of the intensity change could dominate the fit. Therefore, the pixels were weighted by local gradient magnitude to enhance the effect of points near the edge center.

An obvious constraint on the orientation of the line is that it be perpendicular to the gradient of the fitted plane. Thus, this leaves the problem of locating the line along the projection of the gradient. A simple approach is to intersect the fitted plane with a horizontal plane representing the average intensity of the region weighted by local gradient magnitude as shown in Figure 3(b); the straight line resulting from the intersection of the two planes is shown in Figure 3(a).

2.4 Extracting Attributes of the Support Context

The gradient region and the planar fit of the associated intensity surface provides the basic information necessary to quantify a variety of attributes beyond the basic orientation and position parameters. Length is simply the distance between the two endpoints. Other attributes of the line include properties of the intensity profile perpendicular to the line and its behavior along the length of the line. Analysis of the profile of the line can provide a measure of the edge's contrast and width (fuzziness), while behavior along the length determines it's straightness; see [BUR84].

3.0 EXPERIMENTAL RESULTS

The algorithm described in the preceding sections was applied to the full images shown in Figure 1. The algorithm utilized overlapping partitions as described in Section 2.2; the partition size was 45 degrees, staggered by 22.5 degrees. Figures 4-5 demonstrate the performance of the algorithm.

Figure 4(a) shows the unfiltered output of the algorithm applied to the first house image. Note that all of the small and low contrast edges are still present. Figures 4(b-c) show the result of filtering 4(a) on the basis of gradient steepness (change in gray-levels per pixel) followed by a filtering on length that separates the edges into two disjoint sets, one corresponding to short texture edges (Figure 4b) and the other to longer lines related to the surface structure of objects in the image (Figure 4c). We are also examining ways in which texture descriptors may be constructed from the edge set remaining when a filter similar to that which produced 4(b) is applied to the initial edge data. In Figure 4c, the structural edges representing the telephone wires were extracted from a thin, one pixel wide diagonally oriented region, a difficult problem for many line extraction processes.

Figure 5 shows results on another typical house image, this time filtered on length alone (all edges greater than 5 pixels).

4.0 CONCLUSIONS

This paper has presented a low-level representation for straight image lines. The technique for extracting straight lines is effective because it globally organizes the spatial extent of a straight line without local decisions about the meaningfulness of an edge feature. It does this by utilizing gradient orientation to provide a gradient segmentation of the pixels in the formation of edge-support regions. Analysis of the intensity surface of the pixels in these regions yields the information required to extract lines and characterize the intensity variations in a variety of ways. The algorithm is very robust and accurately extracts many low contrast long lines.

While the extracted lines, such as long straight lines, might be directly useful, the underlying edge-support region has a wealth of information useful to intermediate processing strategies. These include additional grouping mechanisms for linking co-linear straight line segments and linking piecewise-linear approximations to curved lines bounding areas of similar properties. It also contains information useful for grouping lines with common properties into textured regions. The representation can serve as a simple yet rich edge-line "primal sketch" [MAR77]. The edge-support regions might be useful in separating the straight lines into intrinsic images [BAR78] representing boundaries of different types such as illumination, texture, reflectance, orientation, etc.

5.0 REFERENCES

[BAR78] Barrow, H.G. and Tenenbaum, J.M., "Recovering Intrinsic Scene Characteristics from Images," in Computer Vision Systems (A. Hanson and E. Riseman, eds.), Academic Press, 1978, pp. 3-26.

[BUR84] Burns, J., Hanson, A. and Riseman, E., "Extracting Straight Lines," COINS Technical Report, University of Massachusetts, January 1984.

[EHR78] Ehrich, R.W. and Foith, J.P., "Topology and Semantics of Intensity Arrays," in Computer Vision Systems (A. Hanson and E. Riseman, eds.), Academic Press, 1978, pp. 111-127.

[HAN80] Hanson, A., Riseman, E., and Glazer, F., "Edge Relaxation and Boundary Continuity," COINS Technical Report 80-11, University of Massachusetts, May 1980.

[HAR81] Haralick, R.M. and Watson, L., "A Facet Model for Image Data," Computer Graphics and Image Processing, Volume 15, 1981, pp 113-129.

[MAR77] Marr, D. and Nishihara, H.K., "Representation and Recognition of the Spatial Organization of Three-Dimensional Shapes," Proc. Royal Society B, 200, 1977, pp. 269-294.

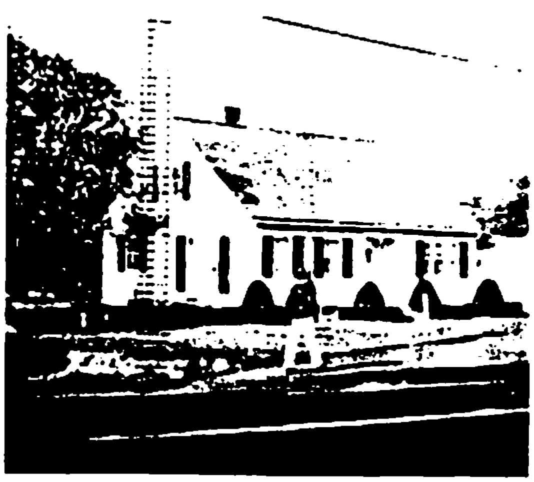

Figure 1. Two images used to demonstrate the straight line finder.

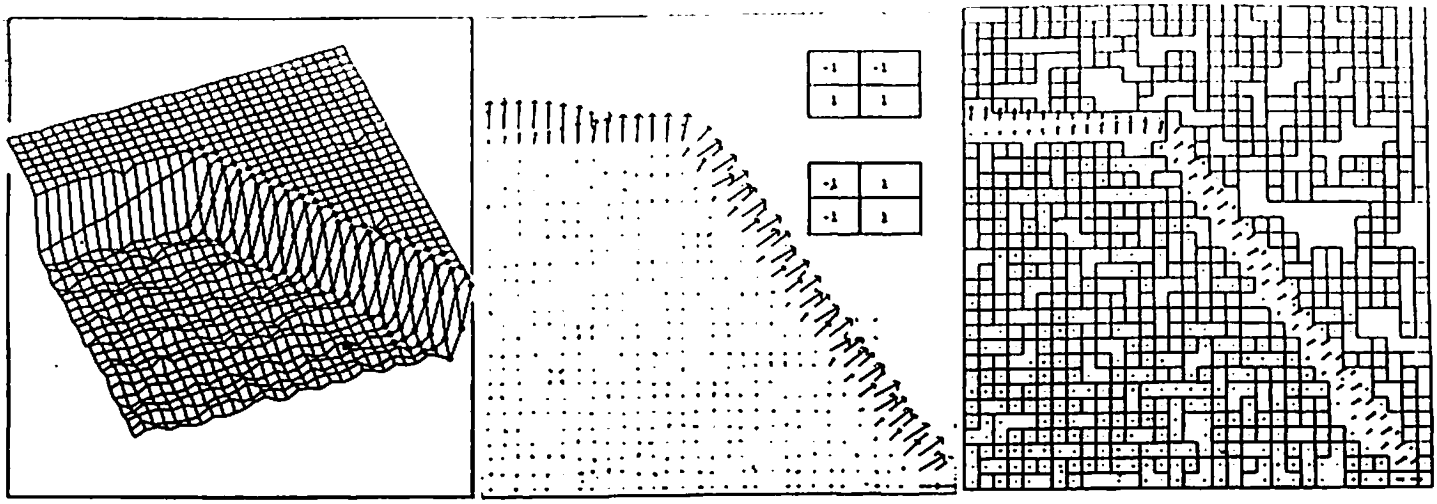

Figure 2. Forming gradient regions. (a) A 32x32 subarea of a house image which will be used to illustrate the process. (b) The 2x2 operators used to estimate dI/dx and dI/dy, from which the local gradient orientation is obtained and the resulting gradient vectors. (c) Gradient regions formed by a regular partitioning of the data into eight orientation classes.

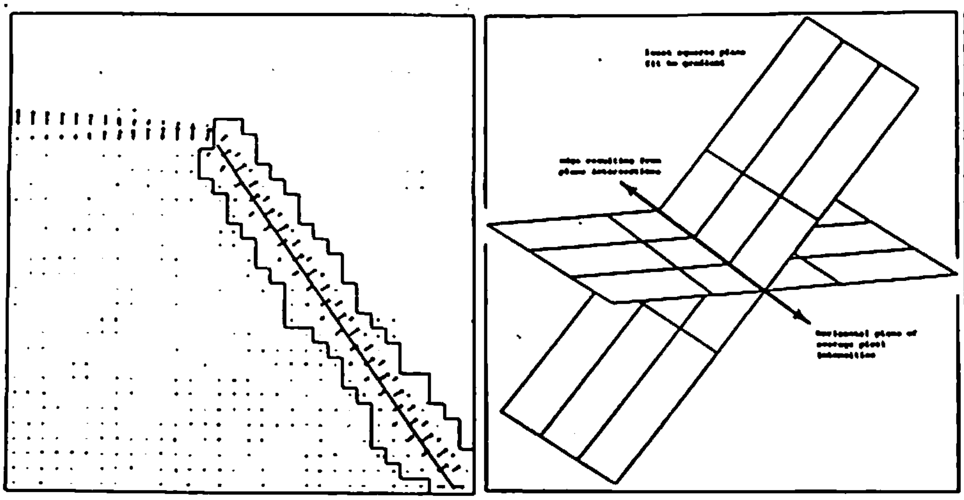

Figure 3. Obtaining straight lines. (a) A group of pixels associated with a gradient region from figure 2 and its interpreted straight line. (b) The straight line is obtained by intersecting the L.S.E. plane fit to the intensities, weighted by gradient magnitudes, with the horizontal plane representing the average intensity, weighted again by magnitude.

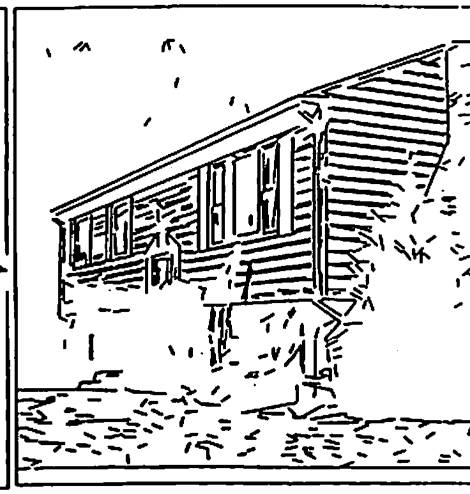

Figure 5. Results of the line extraction algorithm on the house image of Figure 1(b), after filtering on length alone (≥ 5).

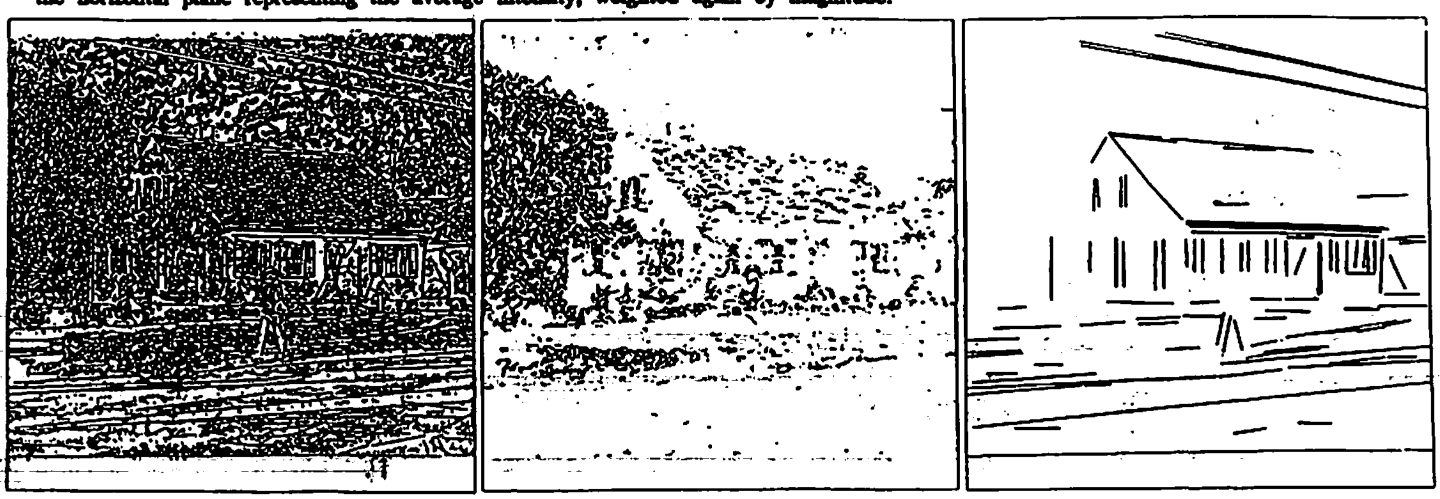

Figure 4. Results of the line extraction algorithm on the house image of Figure 1(a). (a) Initial edge set. (b) Filtering (a) on gradient steepness (≥ 10 gray levels per pixel) and filtering for short edges (length ≤ 5). (c) Filtering for steep long lines (length ≥ 15 and gradient steepness ≥ 10 gray levels per pixel).