EDGE AND LINE DETECTION AS EXERCISES IN HYPOTHESIS TESTING

G.N. Newsam

Intelligence, Surveillance and Reconnaissance Division,
Defence Science and Technology Organisation,
P.O. Box 1500, Edinburgh SA 5111,
AUSTRALIA.
email: garry.newsam@dsto.defence.gov.au

ABSTRACT

In the standard paradigm for edge or line detection matched filters are used to test for the presence of a specified structure. Longer lines are then produced by linking detections; broader lines by filtering at larger scales. The paper proposes an alternative approach based on characterising general linear features as lines along which the distribution of pixel values and/or pixel differences is significantly different from their distribution in the image as a whole. Detection then becomes an exercise in testing the hypothesis that the two distributions are different. For Gaussian distributions the test reduces to computing one or more Radon transforms. Issues of length and scale are now addressed by multiresolution methods: multiresolution implementations of the Radon transform naturally construct longer lines as unions of shorter ones, while broader lines are detected by applying the detection process to the coefficients in each level of a wavelet decomposition of the image.

1. INTRODUCTION

The ideas presented here were spurred by the problem of detecting faint roads and tracks in synthetic aperture radar (SAR) images of outback Australia such as Figure 1. Tracks are characterised by being narrow (often just one pixel wide), darker than the surrounding bush (due to the lower radar reflectivity of the smoother surfaces) and having long straight sections interspersed with wide bends (although the straights rarely extend across the entire image). They can also be quite faint. Standard edge detection algorithms performed poorly on this problem, in part due to the large size of the image and the faintness of the tracks: we therefore sought a different approach.

The dominant paradigm for a formal model of edge or line detection was first formulated by Canny [2]: it is essentially a test of the hypothesis that a given shape (*e.g.* a step edge of a given size) is embedded in the image at a given location. These hypotheses can be stated as

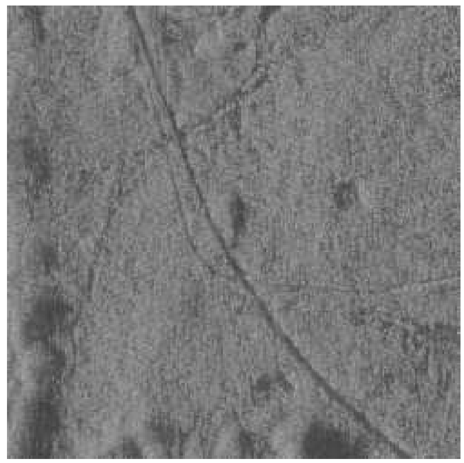

Fig. 1. A SAR image of the Australian bush.

$$\begin{aligned} H_0 : \quad z_{mn} &= w_{mn} \\ H_1 : \quad z_{mn} &= \gamma s_{m-p,n-q} + w_{mn} \end{aligned} \qquad (1)$$

Here z_{mn} is the pixel value at location (m, n) in the image, w_{mn} a random process model of the background (usually assumed to be white noise), s the hypothesised shape, γ the unknown strength of the edge) and (p, q) the postulated location for the shape. The hypothesis test then reduces [9] to thresholding the test statistic

$$t = \frac{(\mathbf{s}_{*-p,*-q}^T \mathbf{z})^2}{\mathbf{s}^T \mathbf{s}\, \mathbf{z}^T \mathbf{z} - (\mathbf{s}_{*-p,*-q}^T \mathbf{z})^2}$$

In practice difficulties with rotation and scale mean that most edge (or line) detection is carried out by thresholding the output of local filters such as Roberts or Sobel operators to identify edge pixels, followed by linking of identified pixels (see [7]). Features at larger scales (*i.e.* smooth edges or broad lines) are found by first convolving the image with a smoothing kernel and downsampling, then repeating the basic detection process.

The above approach has its limitations, both theoretical and computational. For example, edges and lines have different associated ideal shapes s, so require separate tests. Moreover the intuitive definition of an edge or line is broader than that allowed in the paradigm: the boundary between two textures is usually considered to be a line or edge, but it cannot be easily accomodated within the model above. Finally the edge linking and image smoothing processes are not integral to the formal model.

These considerations led to proposal of the following alternative model:

Definition 1 *A* linear feature *in an image is a line along which the distribution of pixel values and/or first differences thereof is significantly different from the distribution of the same quantities in the image as a whole.*

The next section derives the hypothesis test associated with this definition for line detection under the assumption that pixel values on the line and in the image both have Gaussian distributions (albeit with possibly different means and variances). Section 3 describes the extension of this test to edge detection through testing pixel differences along the line as well as values.

The analysis shows that both hypothesis tests require computation of a small number of Radon transforms of the original image and various derived images. Section 4 introduces the new fast multiresolution methods that have been developed for this and notes their advantages. In particular, their structure is such that they compute not only the transform of the entire image but also the transforms on all subimages in a quadratic pyramid. These intermediate results can now be used to detect lines that do not extend across the entire image, and to determine their length.

Scale is always a problem in edge and line detection: lines can have varying widths (and so become parallel edges), while a smooth ramp will not be detectable at small scales. Fortunately, as Section 5 indicates, repeatedly applying the hypothesis test to the various levels of a wavelet decomposition naturally extends the paradigm to encompass scale.

The purpose of this paper is to describe the above new approach; as it is still under development and space here is limited we shall not attempt a description of its performance or a comparison with existing methods. Suffice it to note here that the plan is to implement a multiresolution Radon transform and associated code for line detection within the Analysts' Decision Support System [8], a comprehensive software system being developed by DSTO and partner organisations for operational processing of large volumes of surveillance imagery.

2. LINE DETECTION BY RADON TRANSFORMS

W.l.o.g. histogram equalisation can be used to reduce the distribution of pixel values in the image to a Gaussian with zero mean and unit variance. If in addition the distributions of pixel values along lines are assumed to be independent and Gaussian, then the simplest form of line detection under the new paradigm reduces to testing the hypothesis that the pixel values on a given line have the same univariate Gaussian distribution as those in the image as a whole.

More precisely, suppose that the distribution of the pixel values z_{mn} along the line $\ell(\rho, \theta)$ in the image at a distance ρ from the origin and an angle θ to the x-axis has mean $\mu_{\rho,\theta}$ and variance $\sigma^2_{\rho,\theta}$. We wish to test the hypotheses:

$$\begin{array}{llcl} H_0: & \mu_{\rho,\theta} = 0 & \text{and} & \sigma^2_{\rho,\theta} = 1 \\ H_1: & \mu_{\rho,\theta} \neq 0 & \text{or} & \sigma^2_{\rho,\theta} \neq 1 \end{array} \quad (2)$$

This is a standard problem in statistics (see, *e.g.*, [4]) and is usually settled by the *likelihood ratio test*, *i.e.* by thresholding the statistic

$$\begin{aligned} t(\rho, \theta) &= \log \frac{\prod_{(m,n)\in\ell(\rho,\theta)} pr\left(z_{mn} \in \mathcal{N}(\mu_{\rho,\theta}, \sigma^2_{\rho,\theta}\right)}{\prod_{(m,n)\in\ell(\rho,\theta)} pr\left(z_{mn} \in \mathcal{N}(0,1)\right)} \\ &= \frac{|\ell(\rho,\theta)|}{2}\left[\mu^2_{\rho,\theta} + \sigma^2_{\rho,\theta} - \log \sigma^2_{\rho,\theta} - 1\right] \quad (3) \end{aligned}$$

where $|\ell|$ is the length in pixels of that portion of ℓ in the image.

To compute t it suffices to compute the mean and variance of the pixels along each line. These in turn can be computed given the sums of z_{mn} and of z^2_{mn} along each line. By definition the sums $|\ell(\rho,\theta)|\mu_{\rho,\theta}$ along all lines comprise the *Radon transform* $(\mathcal{R}z)(\rho,\theta)$ of the image (it is also referred to as the *Hough transform* in image processing; see, *e.g.*, [7]). A detailed description of the Radon transform and its properties can be found in [5].

Direct calculation of the Radon transform of an image with $N \times N$ pixels by summation along all lines requires $O(N^3)$ operations, so is prohibitive for large images such as those produced by SAR. Fortunately new algorithms based on multiresolution methods reduce this cost to $O(N^2 \log N)$. Moreover the structure of these algorithms makes possible to detect lines of varying lengths. Further discussion of this issue is postponed to Section 4

Finally, in principle the statistics of t are known and can be used to decide whether or not a specified line is present at any specified significance level. In practice detection is most likely to be done by computing t for all lines in the image and then determining its peaks.

3. EDGE DETECTION

Lines in an image can be viewed as degenerate edges, and edges (and boundaries in general) are usually defined by some degree of discontinuity in pixel values across them. Therefore a general linear feature detector should test both the distribution of pixel values z_{mn} and the distribution of the central differences

$$\begin{aligned} dx_{mn} &\equiv z_{m+1,n} - z_{m-1,n} \\ dy_{mn} &\equiv z_{m,n+1} - z_{m,n-1} \end{aligned}$$

along the line $\ell(\rho, \theta)$. This section develops such a test.

If there is no overall intensity gradient across the entire image, then dx_{mn} and dy_{mn} have zero mean and are uncorrelated with z_{mn}. While not necessary, to simplify the structure of the edge detector we shall assume this holds. We also assume that the differences are independent Gaussian, and that pixel values are uncorrelated with differences along any given line. Thus if the vector $\mathbf{d}_{mn}$ is defined as

$$\mathbf{d}_{mn} \equiv \begin{bmatrix} dx_{mn} \\ dy_{mn} \end{bmatrix}$$

then we assume $\mathbf{d}$ has the Gaussian distribution $\mathcal{N}(\mathbf{0}, \Sigma)$ (for ease of notation from hereon we shall suppress the subscript mn where possible). Here $\Sigma = E[\mathbf{d}\mathbf{d}^T]$ is the (easily computed) covariance matrix for the differences in the whole image.

Now consider the distribution of $\mathbf{d}$ along a given line $\ell(\rho, \theta)$: we assume it is also Gaussian with mean $\mathbf{m}_{\rho,\theta}$ and covariance matrix $C_{\rho,\theta}$. Then edge detection reduces to testing the hypotheses

$$\begin{aligned} H_0 : \quad & \mathbf{m}_{\rho,\theta} = \mathbf{0} \quad \text{and} \quad C_{\rho,\theta} = \Sigma \\ H_1 : \quad & \mathbf{m}_{\rho,\theta} \neq \mathbf{0} \quad \text{or} \quad C_{\rho,\theta} \neq \Sigma \end{aligned} \tag{4}$$

Again following [4] the likelihood ratio test for this is to threshold the statistic

$$\begin{aligned} T(\rho, \theta) \quad = \quad & \frac{|\ell(\rho, \theta)|}{2} \left[\mathbf{m}_{\rho,\theta}^T \Sigma^{-1} \mathbf{m}_{\rho,\theta}^T + \operatorname{tr} C_{\rho,\theta} \Sigma^{-1} \right. \\ & \left. - \log \left| C_{\rho,\theta} \Sigma^{-1} \right| - 2 \right] \end{aligned} \tag{5}$$

Calculating T now requires computing five Radon transforms; the transforms of the images of differences dx and dy, and of the products dx^2, dy^2 and $dx\, dy$.

Under the assumption that pixel values and pixel differences are uncorrelated, the hypothesis tests in (2) and (4) can be combined into a single test for a broader class of linear features by adding together the statistics in (3) and (5).

4. COMPUTING THE RADON TRANSFORM

Until quite recently, the Radon transform has been computed either directly or by Fourier transform methods [5]. In the last few years, however, a new family of methods has been developed based on multiresolution analyses [1, 3, 6]. While it is not possible to provide a full description of these approaches here, we shall sketch out the principle underpinning the algorithms and indicate the useful extra benefits this has for edge detection.

4.1. Multiresolution algorithms

Figure 2 illustrates the basic idea underlying all multiresolution algorithms: the value of the transform on a larger domain can be can be quickly computed given its values on a set of subdomains. The figure shows this in that the integral along the line through the large square is just the sum of the integrals along the three lines formed by the intersection of the original line with the square's four quadrants.

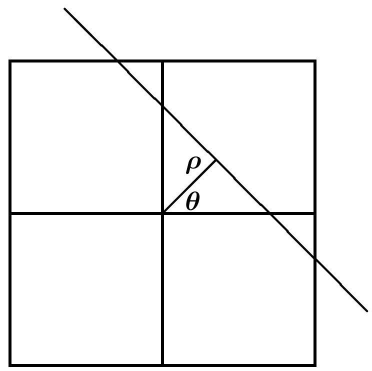

Fig. 2. Multiresolution decomposition of line integrals through a square.

With appropriate care with various approximations, this principle can now be used to construct discretisations of the Radon transform. Given an image z with sides of length $N = 2^K$ pixels, decompose into a standard image pyramid. On each square $\mathcal{S}$ of size 2^k the algorithm then constucts a discretisation containing $O(2^{2k})$ coefficients of $(\mathcal{R}z)(\rho, \theta)$ of z in $\mathcal{S}$ in which each coefficient is written as the sum of corresponding coefficients of the discretised transforms of z in the four squares of size 2^{k-1} that make up $\mathcal{S}$.

The cost of forming these combinations is linear in the number of coefficients, *i.e.* is $O(2^{2k})$. Thus the cost of forming all Radon transforms of squares at level k from those at level k-1 is $O(N^2)$. As there are $O(\log N)$ levels, the cost of computing the Radon transform of the entire image is clearly $O(N^2 \log N)$.

4.2. Determining line length

Unlike Fourier or direct methods for calculating the Radon transform, multiresolution methods calculate not only the integrals along all lines through the entire image but also the integrals along lines in all subimages in the image pyramid. Clearly these intermediate results can be used to test for the presence of lines in the subimages, providing the ability to detect lines that do not extend across the entire image. Moreover we can determine the length of such lines by comparing transform values $(\mathcal{R}z)(\rho, \theta)$ in neighbouring squares for appropriate values of ρ.

We shall not attempt a detailed description of how one might determine the presence and extent of partial lines: we shall only note that this may necessitate computing the transform on overlapping sets of subsquares, and in storing detections in reasonably sophisticated hierarchical data structures. Pointers as to how this can be done efficiently can be found in the literature on multiresolution methods, for example *à trous* wavelet decompostions [10].

5. DETECTIONS ON MULTIPLE SCALES

As noted in the introduction, finding edges or lines at varying scales in the standard paradigm requires repeated convolutions of the images with scaled shapes $s_\alpha(x) \equiv s(\alpha x)$. Increasingly, however, as exemplified in the JPEG2000 standard [11], images are represented in multi-resolution formats rather than by their raw pixel values. Using the algorithms described above to carry out edge or line detection at each level in such a representation is the natural way to search for linear features on a complete range of scales within the image.

To illustrate this, we briefly consider how edge and line detection may be naturally combined when the decomposition is formed from the tensor product of a one dimensional wavelet ψ and its associated scaling function ϕ. In this case the coefficients at each scale in the decomposition can be viewed as vectors $\mathbf{c}_{kmn}$, where

$$\mathbf{c}_{kmn} \equiv \begin{bmatrix} \int\int z(x)\,\phi(2^k(x-m))\,\psi(2^k(y-n))\,dx\,dy \\ \int\int z(x)\,\psi(2^k(x-m))\,\phi(2^k(y-n))\,dx\,dy \\ \int\int z(x)\,\psi(2^k(x-m))\,\psi(2^k(y-n))\,dx\,dy \end{bmatrix}$$

(for simplicity here the underlying image z is assumed to be a a continuous rather than discrete object). Any linear feature, edge or line, at the k-th scale should now manifest itself in the values of $\mathbf{c}_{kmn}$ along the corresponding line. As in Section 3, this hypothesis can be formally tested by comparing the distribution of the vectors along a line with their distribution in the level as a whole under the assumption that both are Gaussian. The resulting test has the same form as that in (5).

6. CONCLUSIONS

We have presented a new paradigm for detecting linear features by testing the hypothesis that the distribution of pixel values, pixel differences or wavelet coefficients along the line is significantly different from their distribution in the image as a whole. In the ideal (*i.e.* independent Gaussian) case, the test requires the calculation of a small number of Radon transforms to determine the necessary means and covariances. These can be computed efficiently via new multiresolution methods, with the added bonus that the algorithms also compute the transforms on all subimages, allowing detection of lines of arbitrary length. Linear features at varying scales can be detected by applying the test to the coefficients on each level of a wavelet decomposition of the image.

7. REFERENCES

[1] A. Brandt and J. Dyn, "Fast calculation of multiple line integrals," *SIAM. J. Sci. Comp.*, **20**(4), pp. 1417–1429, 1999.

[2] J. Canny, "A computational approach to edge detection," *IEEE Trans. PAMI*, **8**(6), pp. 679–698, 1986.

[3] W.A. Götz and H.J. Druckmüller. "A fast digital Radon transform—an efficient means for evaluating the Hough transform." *Pattern Recognition*, **29**(4), pp. 711–718, 1996.

[4] K.V. Mardia, J.T. Kent and J.M. Bibby, *Multivariate Analysis*, Academic Press, London, 1979.

[5] F. Natterer. *The Mathematics of Computerized Tomography*, John Wiley, Chichester, 1986.

[6] G.N. Newsam, J.F. Magarey and T.M. Payne, "A fast hierarchical algorithm for computing the Radon transform of a discrete image," *in preparation*.

[7] W.K. Pratt, *Digital Image Processing*, John Wiley, New York, 2001.

[8] N.J. Redding, "Design of the Analysts' Detection Support System for Broad Area Aerial Surveillance," DSTO-TR-0746 (see www.dsto.defence.gov.au), Australia, 1998

[9] L.L. Scarf, *Statistical Signal Processing: Detection, Estimation and Time-Series Analysis*, Addison-Wesley, Reading, Mass., 1991.

[10] J.-L. Starck, F.D. Murtagh and A. Bijaoui, *Image Processing and Data Analysis: The Multiscale Approach*, Cambridge University Press, Cambridge, 1998.

[11] D.S. Taubman and M.W. Marcellin, *JPEG 2000: Image Compression Fundamentals, Standards and Practice*, Kluwer International Series in Engineering and Computer Science, 2001