Edges: Saliency measures and automatic thresholding

Paul L. Rosin
Institute for Remote Sensing Applications
Joint Research Centre
I-21020 Ispra (VA), Italy
paul.rosin@jrc.it

Abstract — Edges are useful features for structural image analysis, but the output of standard edge detectors must be thresholded to remove the many spurious edges. This paper describes experiments with techniques for: 1) automatically determining appropriate edge threshold values, and 2) alternative edge saliency measures to gradient magnitude.

INTRODUCTION

In recent years edge based techniques from machine vision have been applied to remotely sensed imagery. This can involve delineating regions which can then be processed as complete units. Classifying regions rather than individual pixels has the advantage that their average spectral values are less prone to local fluctuations. Moreover, shape properties can be measured which can aid the classification task. Alternatively, if particular structural objects are to be identified in the scene (e.g. airports) then edges are the most commonly used feature.

One drawback with using edges is that not only do edge detectors extract meaningful and useful edges, but also many other spurious ones which arise from noise and minor changes in spectral values. If all such edges are kept then the resulting clutter is hard for subsequent processing stages to analyse, while the large number of edge points can seriously degrade computational performance. The alternative is to select a subset of edges for further analysis, and disregard the remainder. This is generally done by a threshold on the gradient magnitude of pixels.

AUTOMATIC THRESHOLDING

Unfortunately edge thresholding is normally done in an *ad-hoc* manner, often requiring user tuning of parameters. Standard image intensity thresholding algorithms cannot be applied since they assume bimodal (or multi-modal) intensity histograms while edge magnitude histograms are more likely to be unimodal.

Voorhees and Poggio [12] showed that if the image noise consists of additive Gaussian noise then the magnitude of the gradient of the image

$$\|\nabla I\| = \sqrt{I_x^2 + I_y^2}$$

has a Rayleigh distribution

$$R(z) = \frac{z}{\sigma^2} e^{-\frac{z^2}{2\sigma^2}}.$$

For an acceptable proportion of false edges (P_F) this allows a suitable threshold to be selected:

$$\tau = \sigma\sqrt{-2\ln(P_F)}.$$

However, in practice the edge distribution in a real image is a combination of different sources of noise and significant features, complicating the identification of the non-noise components in the edge magnitude histogram. Voorhees and Poggio's approach was to smooth the histogram and attempt to locate the peak. Unfortunately this requires various heuristically set parameters which diminishes the robustness of the approach. Alternatively, Amodaj and Popovic [1] iteratively fit a Raleigh function to the lower portion of the histogram. However, the fitting may not be robust since it depends on the portion of the histogram used as well as on the initial estimate.

A simpler approach is to estimate the peak directly [4] or iteratively [**?**] from the raw edge magnitude histogram, since $\sigma = \sqrt{\frac{2}{\pi}}\overline{R(z)}$. A potentially more reliable procedure is to use estimation techniques from robust statistics such as the median and least median of squares (LMedS) [10].

Rather than threshold edge pixels independently some contextual information can be incorporated. For instance, Canny [4] employs two thresholds using hysteresis, and a reasonable lower threshold was experimentally determined as half the high one. We have previously described an automatic thresholding technique that operates on complete curves, and does not make assumptions about the noise distribution [11]. It's main dependence is on reasonable connectivity, so that edge curves can be extracted from the image. Unfortunately, although the method worked well on man-made and natural scenes, the poorer definition of objects in satellite imagery caused problems. Possibly this could be rectified using sophisticated edge linkers [5, 9].

SALIENCY MEASURES

Although gradient magnitude is the primary measure used to discriminate between good and bad edges other possibilities have also been suggested.

Lifetime — This involves blurring the image and tracking edges over scale. Tracking can be performed in a fine-to-coarse manner [3], or coarse-to-fine [2]. The longer an edge persists the more likely it is to be significant.

Wiggliness — The expectation is that noisy edges are less likely to be locally straight or smooth than significant edges. One measures of edge wiggliness is angular dispersion [7]. A drawback with such approaches is that they generally require a parameter specifying the window size within which the wiggliness measure is calculated.

Width — The spatial width of an edge may also be useful to discriminate between different edge types. As mentioned by Prager [8], this can be implemented after non-maximal suppression. The strength of each retained edge is found in the unsuppressed edge map by integrating the set of monotonically decreasing gradient magnitudes on both sides of the edge.

Projection onto Edge-Subspace — Frei and Chen [6] suggested that the local 3×3 image window should be projected onto a set of nine orthogonal feature basis functions. Thresholding would then be based on the angle of window's projection onto the edge subspace formed by four of the bases.

EXPERIMENTS

Fig. 1a shows the 1st principal component of a 256×256 portion of a Landsat TM image taken in Portugal. The Canny [4] edge detector ($\sigma = 1$) is independently applied to the six non-thermal bands, and the results ORed together (fig. 1b). Without thresholding (fig. 1c) there are many spurious edges. The most significant edges were selected by hand to provide an approximate reference set for assessing the measures and thresholds (fig. 1d).

The distribution of edge gradient magnitudes only roughly approximates a Rayleigh distribution (fig. 2). It can be seen that the LMedS provides the best estimate of the mode compared to the other normalised estimators, being relatively unbiased by the fat tail which corresponds to non-noisy edges.

The various saliency measures are rated by applying them to the reference edge data and plotting in fig. 3 their ROC (receiver operating characteristic) curves. These demonstrate the trade-offs between the amount of true and false positives for all the different threshold values. In addition, each saliency measure can be rated by the area under the ROC curve. It is clearly seen that the gradient magnitude outperforms all the other salience measures in discriminating between significant and spurious edges.

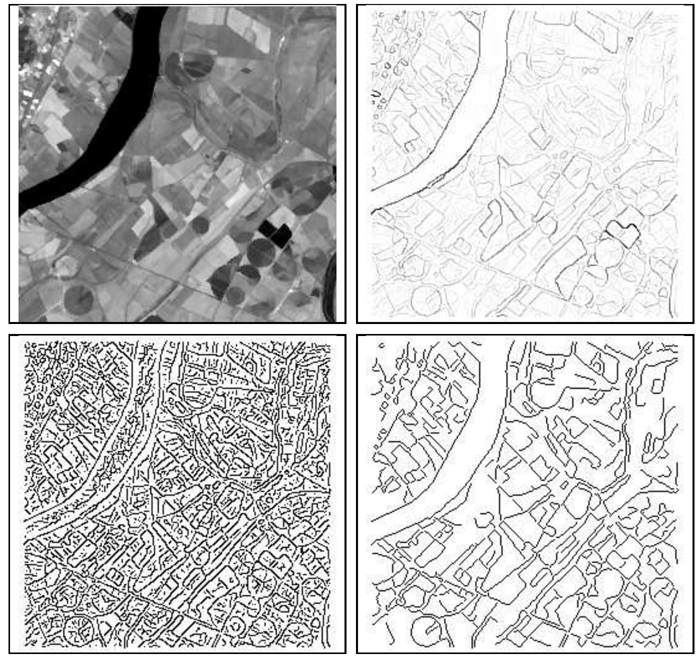

Figure 1: Test image and edges; a) 1st principal component; b) Canny edge magnitudes c) All Canny edges d) Hand-selected Canny edges

This is verified in table 1 which gives for each saliency measure the optimal assessment measure (over all thresholds). The assessment measure is calculated as the product of the proportions of true positives and true negatives [11].

Automatic thresholding was performed on the gradient magnitudes using the LMedS mode estimator with 90% confidence in rejecting noisy edges, determining the threshold shown in fig. 2. However, the result (fig. 5a) appears overthresholded. This is confirmed by the assess-

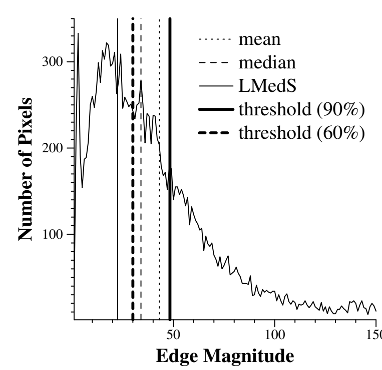

Figure 2: Edge gradient magnitude distribution

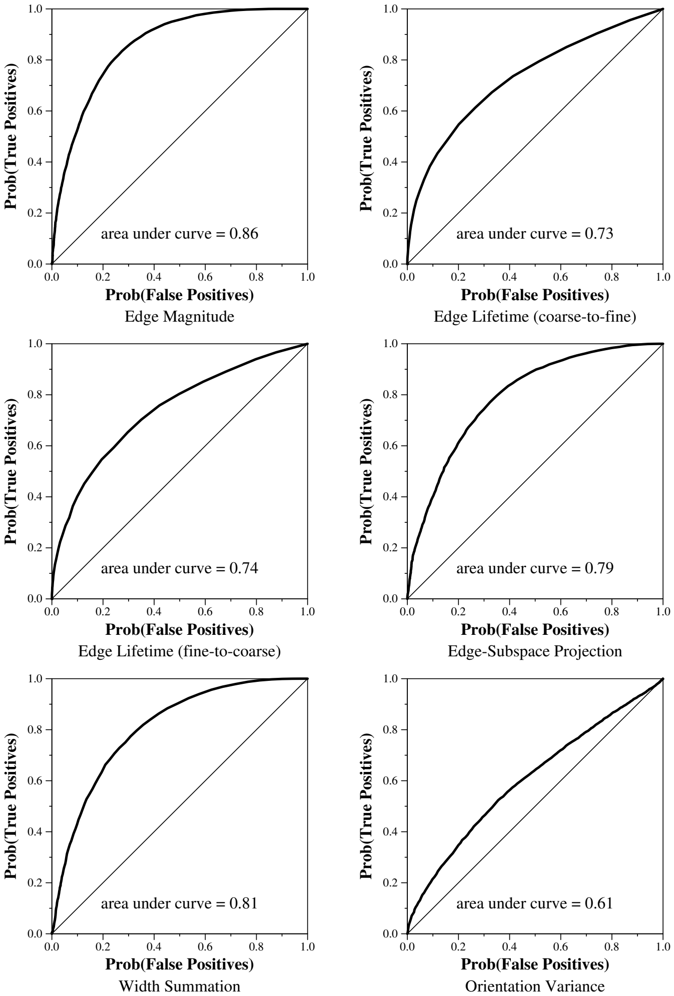

Figure 3: ROC curves of saliency measures

ment measure curve in fig. 4, since the selected threshold (the dotted line at 48) is greater than the optimal threshold at 30 (fig. 5b). This occurs since, compared to machine vision images, remote sensed images have many more significant edges with low magnitudes, requiring a lower threshold determined by $P_F \approx 0.4$.

REFERENCES

[1] N. Amodaj and M. Popovic. Adaptive procedure for threshold selection in directional derivative edge detectors. In *Proc. SPIE Applications of Digital Image Processing XIII Vol. 1349*, pages 128–138, 1990.

[2] F. Bergholm. Edge focusing. *IEEE Trans. PAMI*, 9:726–741, 1987.

[3] W.F. Bischof and T. Caelli. Parsing scale-space and spatial stability analysis. *Computer Vision, Graphics and Image Processing*, 42:192–205, 1988.

[4] J. Canny. A computational approach to edge detection. *IEEE Trans. PAMI*, 8:679–698, 1986.

Table 1: Assessment results of saliency measures

SALIENCY MEASURE	OPTIMAL ASSESSMENT
gradient magnitude	0.607
width summation	0.536
edge projection	0.521
lifetime (fine → coarse)	0.459
lifetime (coarse → fine)	0.450
orientation variance	0.338

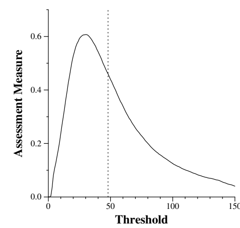

Figure 4: Assessment measure

[5] M.A. Fischler. The perception of linear structure: a generic linker. In *ARPA Image Understanding Workshop*, 1994.

[6] W. Frei and C.C. Chen. Fast boundary detection: a generalization and a new algorithm. *IEEE Comp.*, 26:988–999, 1977.

[7] P.H. Gregson. Using angular dispersion of gradient direction for detecting edge ribbons. *IEEE Trans. PAMI*, 15:682–696, 1993.

[8] J.M. Prager. Extracting and labelling boundary segments in natural scenes. *IEEE Trans. PAMI*, 2:16–27, 1980.

[9] C. Rothwell, J. Mundy, B. Hoffman, and V.D. Nguyen. Driving vision by topology. Technical Report 244, INRIA, 1994.

[10] P. Rousseeuw and A. Leroy. *Robust Regression and Outlier Detection*. Wiley, 1987.

[11] S. Venkatesh and P.L. Rosin. Dynamic threshold determination by local and global edge evaluation. *Computer Vision, Graphics and Image Processing*, 57(2):146–160, 1995.

[12] H. Voorhees and T. Poggio. Detecting textons and texture boundaries in natural images. In *Proc. ICCV*, pages 250–258, 1987.

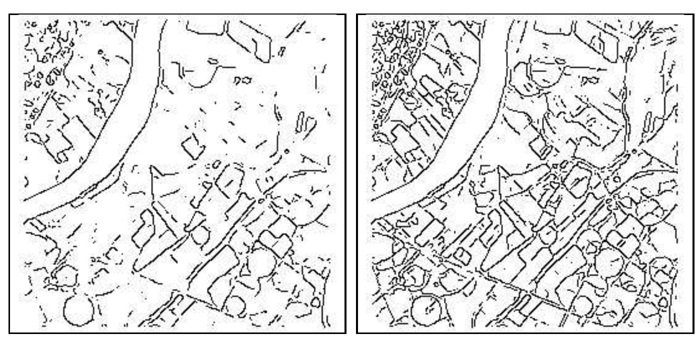

Figure 5: Thresholded edges; a) $P_F = 0.1$; b) $P_F = 0.4$