Classifying Land Development in High Resolution Satellite Images Using Straight Line Statistics

Cem Ünsalan and Kim L. Boyer
Signal Analysis and Machine Perception Laboratory
Department of Electrical Engineering
The Ohio State University
Columbus, OH 43210-1272 USA

Abstract

We introduce a set of measures based on straight lines to assess land development levels in high resolution (1 meter) satellite images. Urban areas exhibit a preponderance of straight line features, generally appearing in fairly simple, quasiperiodic organizations. Wilderness and rural areas produce line structures in more random spatial arrangements. We use this observation to perform an initial triage on the image to restrict the attention of subsequent, more computationally intensive analyses. We extract statistical measures based on straight lines to guide the analysis. We base these measures on orientation, length, contrast, periodicity and location. We trained and tested parametric and non-parametric classifiers using the feature set. Finally, we introduce a decision system performing region classification via an overlapped voting method for consensus discovery.

I. Introduction

A problem of major interest to regional planning organizations, disaster relief agencies, and the military is the identification and tracking of land development across large scale regions, and over time. New, publicly available, high-resolution satellite images offer great promise in attacking these problems. However, the rate at which these new sensors produce data far exceeds the capacity of image analysts to interpret it. To that end, we are developing an autonomous image analysis system to classify land development (for example into urban, suburban, rural, wilderness).

We begin with a set of measures based on straight lines. Subsequent detailed analyses (counting and classifying dwellings, for example) can then be confined to "developed" areas. Straight line structures will be more prevalent and more organized in developed areas than in wilderness or rural areas. Our objective at this stage is the (rough) classification of the image into regions of little or no development, suburban regions likely to be residential, and urban areas.

Our statistical measures are based on the orientation, length, contrast, periodicity, and location of straight lines. We apply a Bayes classifier to label each image region. Initially, we defined a two-class problem to discriminate urban and "not urban" regions and obtained excellent results. Next we addressed the far more difficult problem of discriminating among urban, residential (suburban), and wilderness areas in a three class problem. The direct approach was less successful in this case, largely because suburban regions bridge the other two in our feature space. Therefore, in an attempt to extract suburban regions we introduced two enhancements. First, we defined a transfer function mapping locations in feature space into a measure of "suburbanity." Used alone, this approach was inadequate. But when we introduced a spatial coherence constraint, this time via a novel perceptual grouping approach, the results improved dramatically.

In related work, Lin and Nevatia [1] and Nevatia and Price [2] detected buildings in and other structures in aerial images. Yu *et al.* [3] applied texture analysis to land usage classification. Recently Kim and Muller [4] detected buildings by graph based methods applied to lines extracted by line support regions.

II. Line Support Regions

Burns *et al.* [5] introduced the concept of line support regions, constructed by grouping contiguous pixels of consistent gradient orientation. His gradient calculation used simple 2×2 masks, creating many false alarms in our high resolution data. Therefore, we turned to the scale-controllable IIR edge detection filter introduced by Sarkar and Boyer [6] to compute an optimally smoothed gradient. Then, instead of marking edges in the usual way, we organized the image into line support regions along the lines of Burns' work.

Burns *et al.* [5] obtained his lines by intersecting a horizontal plane with the best fit plane to the local image surface. Following Tan [7] we improve on this approach by fitting an ellipse to each line support region perimeter using a Fourier descriptor approximation. The long axis of the ellipse (which requires only the first order descriptors) defines the straight line. This method is faster and more robust than plane

1051-4651/02 $17.00 © 2002 IEEE

fitting. Three sample subwindows for the wilderness, residential and urban areas are given in Figs. 1, 3, and 5. Lines extracted from these sample sub-windows are given in Figs. 2, 4, and 6, respectively. The relative degree of organization over the three images is evident in their respective line structures.

Figure 1: The wilderness region

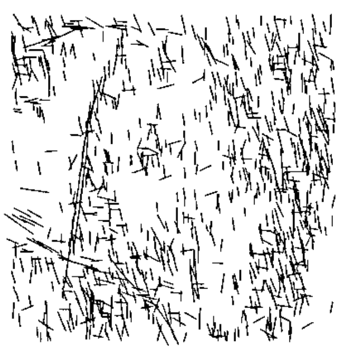

Figure 2: Lines from the wilderness

III. Feature Extraction

We extract features over 400 × 400 pixel windows, with 50% overlap. An organization lying partly in one window will likely lie totally inside a neighboring window. We discard straight lines shorter than five pixels; such lines can't be approximated robustly due to low signal to noise ratios in the contour approximation and certainly don't offer strong evidence for "straightness."

The features are:

- Mean line length, $\bar{\mu}_l$
- Entropy of line length, E_l
- Mean line contrast, $\bar{\mu}_c$
- Entropy of line contrast, E_c
- Orientation distribution, C_o (below)

Figure 3: The residential region

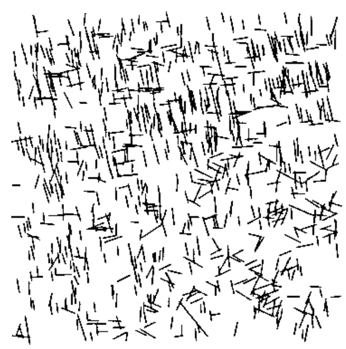

Figure 4: Lines from the residential region

- Entropy of line spacing, E_{dcom} (below)
- Periodicity of line placement, R (below)

The orientation feature is the crosscorrelation between the histogram of line orientations in the window with an idealized bimodal density function model for urban areas, with the modes separated by 90°. To compute the entropy of line spacing, we locate the center of mass of the nearest two lines (to a given line) of similar length (±3 pixels) and build a histogram of these distances. For periodicity of line placement, we project the line segments in the window onto the four axes $y = 0$, $y = x$, $x = 0$, and $y = -x$. In each projection, the line receives weight according to its length. The resulting functions are analyzed using the *periodicity transform* introduced by Sethares and Staley [8], and the maximum retained.

IV. Experimental Results

In this section we present the results of several classification approaches using the features described above, as well as spatial coherence. Our dataset consisting of 28 Ikonos (1-meter) images, 18 drawn from across North America, plus the 10 in Table 1, encompasses a wide range of cultures and development patterns.

Figure 5: The urban region

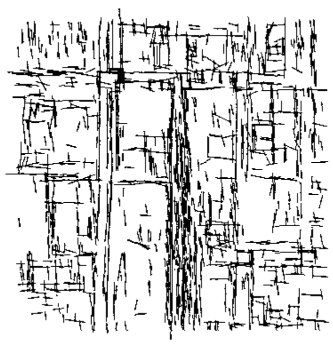

Figure 6: Lines from the urban region

GEOGRAPHIC LOCATION	IMAGE SIZE
Dubai, Arab Emirates	763x953
Sydney, Australia	2003x2003
Rome, Italy	1993x1993
Sanaa, Yemen	1841x2109
Taipei, Taiwan	2101x2101
Tokyo, Japan	1989x2010
Guayaquil, Ecuador	1999x1999
Great Pyramids, Egypt	2104x2104
London, England	2000x2000
Sydney II, Australia	1430x1539

Table 1: Non-North American test images

A. *Two Class Results: Bayes classifier*

An exhaustive search over the full feature set (all singles, pairs, triples, etc.) revealed that the 2D space of contrast mean and entropy offered the best performance. For training we used 48 rural and 48 urban image windows. A *caveat* on ground truth: Some windows straddle transitions from developed to undeveloped areas, for instance, and others encompass lightly developed regions, parkland, and other types of terrain that are hard to classify, even manually. We tested 2014 new region samples and tabulate the results for the best feature space in Table 2.

	percentage classified as	
	wilderness	urban
905 wilderness samples	90.6077	9.3923
1109 urban samples	16.4112	83.5888

Table 2: Bayes confusion matrix, $(\bar{\mu_c}, E_c)$ space

B. *Suburbia and Spatial Coherence*

To this point we have treated each window as a distinct entity without regard to its neighbors. Yet levels of development in neighboring areas are not statistically independent. Cities tend to grow from a core outwards; apart from the intervention of natural barriers such as water and mountains, land development expands more or less coherently. To exploit this observation, as we attempt to recover suburban regions, we have developed a grouping and classification method motivated by principles from perceptual organization.

To define the suburban, or residential class, we begin by considering the distribution of training data depicted in Fig. 7. Clearly, the residential areas, in which we are most interested, bridge the urban and rural areas in feature space, much as they do on the ground. This makes the direct detection of such areas difficult, but we have been able to make significant inroads on the problem that at least suffice for our "triage." We first consider the classification of individual windows (into "suburban" or not), and then introduce the notion of spatial coherence and show how that results in significant improvement.

We begin by estimating the mean vector $\bar{\mathbf{x}}$ and covariance matrix $K_{\mathbf{x}}$ of the suburban class in feature space assuming a bivariate Gaussian distribution. We compute the Mahalanobis distance between each window to be classified and this distribution:

$$D_M = (\mathbf{x} - \bar{\mathbf{x}})K_{\mathbf{x}}^{-1}(\mathbf{x} - \bar{\mathbf{x}})^T \qquad (1)$$

The degree of membership of a given window (or feature vector $\mathbf{x}$) is then computed as (The value 20 was selected experimentally.):

$$\mathcal{M} = e^{-\left(\frac{D_M}{20}\right)} \qquad (2)$$

To test this statistic, we used a Neyman-Pearson decision rule, yielding the receiver operating characteristic (ROC) labeled "Before grouping" in Fig. 8. Although cast as a two class problem (suburbs, or not), this is really a three class issue with fairly unequal priors, and so the ROC – while not spectacular – is far better than chance. It also makes no use (yet) of spatial coherence.

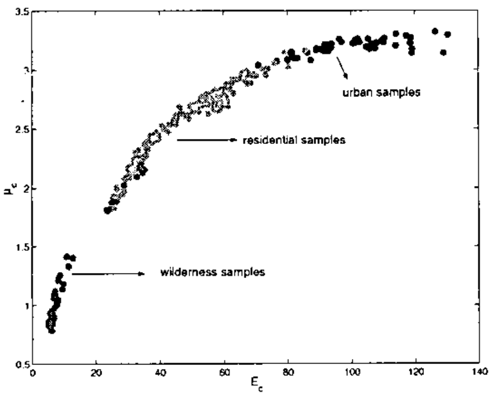

Figure 7: Sample distribution in $\mu_c \& E_c$

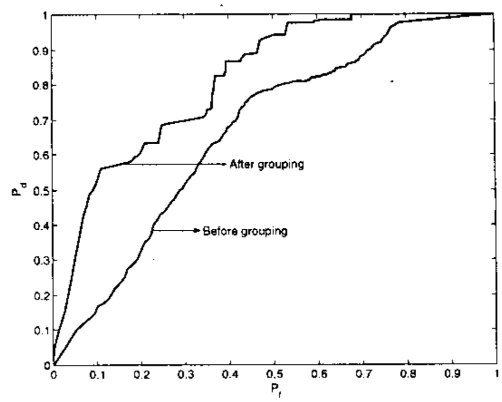

Figure 8: ROC curves for suburbia

Now, to exploit spatial coherence, we frame the problem as one of perceptual grouping. Rather than making hard window by window decisions, we group together windows having compatible (similar) feature values into contiguous regions, and classify them together. In this way we discover *natural groupings* of areas having similar developmental characteristics, as measured in the $(\bar{\mu}_c, E_c)$ feature space.

In grouping, we first construct a 2D "histogram" of the feature vectors over the image. This is not a histogram in the usual sense, but one in which each window votes into all bins satisfying a *compatibility* constraint. This compatibility constraint takes the form of a 2D footprint in feature space; any window having feature values in the range covered by this footprint are considered to be compatible with the current (voting) window in that they have similar features. We then interrogate the histogram to extract contiguous regions having compatible features, from largest to smallest. The extracted collections of windows are called *segments*.

In these experiments, the histograms are of 16 bins each, uniform over the observed range of the corresponding feature value over the training data. The compatibility constraint is three bins, roughly 20% of the range.

We then classify a segment using the membership $\mathcal{M}$ as above. However, a segment's (net) feature vector is that corner of its range corresponding to each feature at its maximum. Although one might expect the average or median value of the feature range to be more representative (as we did), the "upper right" corner value performs better. The reasons are as yet unclear. All windows in the segment then receive the same degree of membership and classification. Again applying the Neyman-Pearson decision rule produces the "After grouping" ROC curve in Fig. 8, which shows substantial improvement.

V. Conclusion

In this paper we present the first stage of a system to assess the degree of land development, with particular emphasis on the extraction of suburban (residential) areas. We base the assessment on the photometric and geometric characteristics of straight line segments. The discrimination of urban and relatively undeveloped areas works very well; the extraction of suburban areas is more challenging but the results are quite encouraging.

VI. References

[1] C. Lin and R. Nevatia, "Building detection and description from a single intensity image," *Computer Vision and Image Understanding*, vol. 72, no. 2, pp. 101–121, November 1998.

[2] R. Nevatia and Price K. E., "Locating structures in aerial images," *IEEE Trans. on PAMI*, vol. 4, no. 5, pp. 476–484, September 1982.

[3] S. Yu, M. Berthod, and G. Giraudon, "Toward robust analysis of satellite images using map information:application to urban area detection," *IEEE Trans on GeoRS*, vol. 37, no. 4, pp. 1925–1939, July 1999.

[4] T. Kim and J.P. Muller, "Development of a graph based approach for building detection," *Image and Vision Computing*, vol. 17, no. 1, pp. 3–14, January 1999.

[5] J.B. Burns, A.R. Hanson, and E.M. Riseman, "Extracting straight lines," *IEEE Trans. on PAMI*, vol. 8, no. 4, pp. 425–455, July 1986.

[6] S. Sarkar and K.L. Boyer, "On optimal infinite impulse response edge detection filters," *IEEE Trans. on PAMI*, vol. 13, no. 11, pp. 1154–1171, November 1991.

[7] S.G. Tan, *Image feature extraction:Line detection and organization*, The Ohio State University, M.Sc. Thesis, Autumn 1990.

[8] W.A. Sethares and T.W. Staley, "Periodicity transforms," *IEEE Trans. on PAMI*, vol. 47, no. 11, pp. 2953–2964, November 1999.