DEVELOPING SPATIAL RE-CLASSIFICATION TECHNIQUES FOR IMPROVED LAND-USE MONITORING USING HIGH SPATIAL RESOLUTION IMAGES

M.J. Barnsley and S.L. Barr

Remote Sensing Unit, Department of Geography,
University College London, 26, Bedford Way,
LONDON, WC1H OAP, England

Abstract. The reasons for the poor performance of conventional, per-pixel classification algorithms applied to satellite sensor images of urban areas are examined. It is argued that standard algorithms are poorly adapted to distinguish between different urban *land-use* categories, particularly in high spatial resolution images, due to the complex spatial pattern of spectrally distinct *land-cover* types in urban areas. Alternative techniques need to be developed which make use of both spectral and spatial information within the scene. This study examines one such technique that attempts to derive information on land use in two stages. Firstly, by performing a low-level segmentation of the image into a few, broad *land-cover* types. Secondly, by grouping the classified pixels into discrete *land-use* categories on the basis of the frequency and spatial arrangement of the class labels. The second stage is performed using a procedure developed in this study, referred to as a SPAtial Re-classification Kernel (SPARK). This examines the number of occasions on which different types of land cover are adjacent to one another, on a pixel-by-pixel basis. This information is used to construct an 'adjacency vector' for the central pixel in the kernel. Land use is inferred by comparing the derived adjacency vectors with those of previously selected sample areas of the candidate land-use categories. Using this technique, an overall accuracy of greater than 85% is obtained for a SPOT-HRV multispectral sub-scene of London.

Key Words: Spatial Re-Classification, Land Use, Land Cover, Urban Areas.

INTRODUCTION

It is widely reported that the results obtained from conventional, multispectral classification algorithms applied to satellite sensor images of urban areas tend to be much less satisfactory than those obtained from equivalent images of agricultural scenes (Forster 1985, Barnsley *et al.* 1989, Sadler and Barnsley 1990). Initially, this was attributed to the relatively coarse spatial resolution of early satellite sensors (Jackson *et al.* 1980, Forster *et al.* 1980). However, despite considerable improvements in sensor spatial resolution in recent years, the expected increases in classification accuracy have rarely been forthcoming (Forster 1985). Indeed, several studies report reductions in classification accuracy using higher resolution data sets (Haack *et al.* 1987, Martin *et al.* 1988). This is usually explained in terms of an increase in '*scene noise*'. In other words, as the spatial resolution increases, discrete scene elements (e.g. buildings, roads and open spaces) begin to dominate the detected response of individual pixels; as a result, the spectral response of urban areas becomes more heterogeneous, making consistent classification problematic (Gastellu-Etchegorry 1990).

Although it is tempting to see this as a problem of sensor spatial resolution, it is perhaps more accurately expressed in terms of the limitations of standard, *per-pixel* classification algorithms and the manner in which they are employed. In pursuing this point, it is worth noting that, while information on the distribution and extent of different *land cover* types (e.g. pavements, roads, roofs, grass lawns, bare soil) within urban areas is important, data on *land use* (i.e. residential, industrial, open space) are of much wider relevance, especially to the town planning community. However, while there is often a relatively simple, direct relationship between land cover and the spectral response detected by a satellite sensor, this is seldom true for land use, particularly in urban areas (Gastellu-Etchegorry 1990, Gong and Howarth 1990, Barnsley *et al.* 1991). Thus, the fundamental problem involved in producing accurate maps of ***land use*** for urban areas is that individual categories of land use frequently represent complex spatial assemblages of a disparate set of ***land cover*** types - each of which may have different spectral reflectance properties (Gong and Howarth 1990, Barnsley *et al.* 1991).

Unfortunately, per-pixel classification algorithms are poorly adapted to deal with this type of spatial variability. This is because they assign each pixel to one of the candidate classes solely on the basis of its spectral properties; no account is taken of the pixel's location within the image or its relationship to the spectral response of neighbouring pixels. Similarly, in the context of a supervised classification, it is extremely difficult to define suitable training areas for many categories of urban land use, due to the variation in the spectral response of their component land-cover types (Forster 1985, Gong and Howarth 1990, Barnsley *et al.* 1991). Thus, training statistics that are derived from contiguous blocks of pixels commonly exhibit both a multi-modal grey-level distribution and a large standard deviation in each spectral waveband (Sadler *et al.* 1991). The implication of the former is that the training statistics for urban areas violate one of the basic assumptions of the widely used Bayesian maximum-likelihood classification algorithm (i.e. that the pixel values follow a multi-variate Normal distribution). The effect of latter is frequently to produce a pronounced overlap between urban and non-urban classes in the multispectral feature space. Finally, these effects may be compounded by the fact that the *mean* grey-levels for urban land-use categories often only differ from those of the non-urban classes in an arbitrary and unpredictable manner, dependent on the location of the training areas (Sadler *et al.* 1991, Barnsley *et al.* 1991).

ALTERNATIVE PROCEDURES FOR INFERRING URBAN LAND USE IN SATELLITE SENSOR IMAGES

If per-pixel classification algorithms are often inappropriate to determine land use directly from satellite sensor images, alternative procedures must be found. Potentially, a more fruitful approach is to develop techniques that make use of the spatial variation in the detected spectral response of urban areas. Various techniques have been suggested for this purpose; these can be divided into three broad groups, namely:

(i) pre-classification image transformations and feature-extraction techniques, such as median filters (Atkinson *et al.* 1985, Sadler *et al.* 1991) and various measures of image texture (Haralick 1979, Franklin and Peddle 1990, Gong and Howarth 1990, Sadler *et al.* 1991)

(ii) contextual classification algorithms (Gurney 1981, Gurney and Townshend 1983, Gong and Howarth 1989)

(iii) post-classification spatial processing, including simple majority filters and spatial re-classification procedures (Wharton 1982, Whitehouse 1990, Guo and Moore 1991).

It is this last group of techniques that will be considered here. In particular, the limitations of simple, frequency-based spatial re-classification algorithms are explored. In view of the shortcomings of these techniques, a new spatial re-classification procedure (SPARK) is developed. This examines the spatial arrangement, as well as the absolute frequency, of different land-cover types within a kernel convolved with a previously classified image. This is used, in turn, to infer the land use associated with every pixel in the image.

Spatial Re-Classification Procedures

The approach adopted in spatial re-classification procedures is to examine the spatial arrangement of different land-cover types within an image. The underlying assumption is that individual categories of land use have characteristic mixtures of spectrally distinct land-cover types that enable their recognition in high spatial resolution images (Wharton 1982). For example, residential districts in many western European cities are characterized by the intermixing of roofs, roads and gardens. Implementation of this approach involves an initial, low-level segmentation of the image into separate land-cover types. The spatial arrangement of these class labels can then used to assign each pixel to a specific land use (Wharton 1982). This approach will be referred to as spatial or contextual *re*-classification.

An example of frequency-based spatial re-classification is provided by Wharton (1982), who performed an unsupervised segmentation of land cover and then calculated the frequency of different cover types within a 3x3 pixel kernel convolved with the classified image. Land-use categories were derived using an unsupervised, non-parametric clustering procedure applied to the frequency data. Similar techniques have recently been employed by Whitehouse (1990) and Guo and Moore (1991); although, in these studies, the frequency distribution of land-cover types surrounding each pixel was compared with those of known areas of the candidate land-use categories.

Although Wharton's method examines the frequency with which different class labels occur within the kernel, no account is taken of their spatial arrangement. The limitation of this approach is evident from the following example. Consider two separate 3x3 pixel windows, each of which has four pixels labelled as the land-cover class 'Built'. In an industrial or commercial area, where these might represent a single large building, the pixels are likely to be clustered together in a block (Figure 1a); in a residential area, where the same class labels might represent individual houses, the 'Built' pixels may be arranged in a line (terraced housing) or may be physically separate (detached housing) (Figure 1b). However, a procedure which simply calculates the frequency of different class labels within the window has no means of distinguishing between these two situations.

THE SPATIAL RE-CLASSIFICATION KERNEL

The example given in the previous section illustrates the need to find a simple method for recording both the frequency and the spatial arrangement of class labels within any given region of an image. One way of doing this is to record the number of times that different class labels occur next to one another within a pre-defined, moving window. A simple technique to achieve this, referred to as the SPAtial Re-classification Kernel (SPARK), has been developed in this study.

Description and Operation of SPARK

SPARK operates by examining all possible pairs of adjacent pixels within a square kernel (i.e. those connected along an edge or by a vertex; Figure 2); the size of the kernel is selected by the user. The class label associated with each of the pixels is noted. This is used to determine the frequency with which different classes are adjacent to one another within the kernel. Thus, in Figure 1a, there are six occasions on which 'Built' pixels are adjacent to one another (hereafter, this will be referred to as a Built-Built adjacency). In the same window, there are also four occurrences of Built-Tree adjacency, and so on. Although the kernel in Figure 1b contains exactly the same number of pixels belonging to each class, there are only three occurrences of Built-Built adjacency, but six of Built-Tree. This demonstrates that, unlike the method developed by Wharton (1982), SPARK is sensitive to the spatial arrangement of land-cover types, as well as their absolute frequency. Therefore, SPARK may be expected to distinguish between somewhat more subtle differences in land use.

a)

G	B	B
T	B	B
G	T	G

Key:- B = Building
G = Grass
T = Tree

b)

B	G	G
T	B	G
B	T	B

Figure 1 Simulated 3x3 pixel window for a) an industrial/commercial area, b) a residential area.

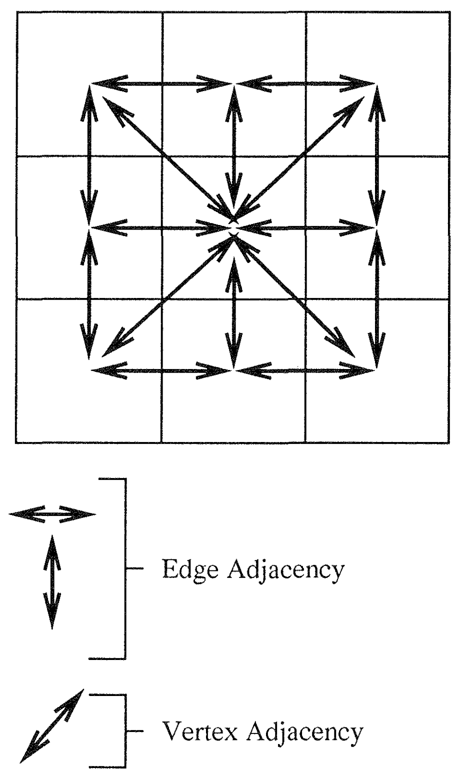

Figure 2 Diagram showing connections between pixels in a 3x3 kernel

In practice, SPARK is convolved with the land-cover image to produce an *adjacency vector*, A, for each pixel:

$$A_{i,j}\ (f_{1,1},\ f_{1,2},\ f_{1,3},\ldots\ f_{x,y}) \tag{1}$$

where the subscripts i and j denote the position of the pixel within the image. The value of each element, $f_{x,y}$, of the vector denotes the frequency with which pixels belonging to class x are adjacent to those belonging to class y for the current position of the kernel. The length, L, of the adjacency vector is conditioned by the number of classes, C, in the image:

$$L = \sum_{m=0}^{C} (C-m) \tag{2}$$

Thus, for an eight class image, each vector will contain 36 elements. For computational purposes, the adjacency vector may be thought of as an array of integers. For most studies, where the number of land-cover classes is reasonably small, this represents an efficient means of storing information about the spatial arrangement of land-cover types within the image.

<u>Inferring Land-Use by Comparison of Adjacency Vectors with 'Template' Vectors</u>

The *land-use* category of a given pixel is determined by comparing its adjacency vector with those derived from representative sample areas of the candidate land uses; the latter will be referred to as *'template'* vectors. The size of the sample areas used to generate each template vector is the same as that of the kernel used in the spatial re-classification stage (i.e. 3x3, 5x5, 7x7 pixels etc). Multiple templates vectors may be defined for each land use. These may either be used independently, or pooled to produce an 'average' template vector. The advantage of using a series of independent templates for a single land use is that subtle variations in the spatial arrangement of its constituent land-cover types at different locations within the image can be taken into account. However, it results in a linear increase in computation time.

As the kernel is passed over the image, the current adjacency vector is compared with each of the land-use templates using the following equation:

$$\Delta_k = \frac{1}{2(N^2)}\sum_{n=0}^{L} \{A_{i,j}(n) - T_k(n)\}^2 \tag{3}$$

$$0 \leq \Delta_k \leq 1 \tag{4}$$

where

Δ_k	=	a measure of similarity between the current adjacency vector and the land-use templates,
$A_{i,j}(n)$	=	element *n* of the current adjacency vector,
$T_k(n)$	=	element *n* of the template vector for land-use category *k*,
N	=	the total number of adjacency events in the kernel (This is determined by the size of the kernel; e.g. N=20 for a 3x3 pixel kernel).

Equation (3) examines the difference between the value of each element of the current adjacency vector, $A_{i,j}(n)$, and the value of the corresponding element in one of the template vectors, $T_k(n)$. The difference is then squared to remove negative values and to highlight large deviations between the two vectors. The sum of the squared differences is divided by a factor, $2(N^2)$, to normalize the result with respect to the size of kernel. This scales Δ_k to a range of 0 to 1; where a value of 0 indicates a perfect match between the current adjacency vector and one of the land-use templates (i.e. no difference between the two vectors), and a value of 1 indicates no match.

Equation (3) has been used by Barnsley *et al.* (1991) in initial tests of SPARK. However, recent experiments suggest that this equation is relatively insensitive to small differences in the frequency and arrangement of the class labels. Therefore, in this paper, we propose a modification to the original equation, namely application of a square root transformation:

$$\Delta_k = \sqrt{\frac{1}{2(N^2)}\sum_{n=0}^{L} \{A_{i,j}(n) - T_k(n)\}^2} \tag{5}$$

The centre pixel in the kernel is assigned to the land-use category for which Δ_k is minimized:

$$P_{i,j} \leftarrow k \ where\ \Delta_k = \min(\Delta_1, \Delta_2, \ldots, \Delta_t) \leq \delta \tag{6}$$

where

$P_{i,j}$	=	the pixel corresponding to the centre of the kernel,
t	=	the number of individual templates,
δ	=	a user-specified threshold.

A user-specified threshold, δ, can be set to prevent pixels being assigned to a land-use category on the basis of a weak match (i.e. Δ_k close to 1.0). This was not used in the present study.

One of the strengths of this procedure is that the memory requirements are very small. This is because, at any given moment in time, the program only needs to hold the adjacency vector for the current location of the kernel and those of the land-use templates.

STUDY AREA AND SATELLITE SENSOR DATA

For the purpose of this investigation we have concentrated on an area to the south-east of London, covering the Borough of Bromley. This area encompasses various different types of land use, ranging from densely-occupied early 20th century housing in the north-west, through major shopping areas and inter-war industrial areas in the centre, to low-density suburbs in the south-east. Surrounding the urbanised area are very large tracts of open country, much of which is statutorily protected Green Belt land.

The data used in this investigation have been extracted from a cloud-free, multispectral (XS) SPOT-HRV image of London (scene 32, 246; +22.46°) acquired on 30th June 1986. In particular, a 512x512 pixel sub-section (approximately 10km x 10km) of the full image, centred on the town of Orpington, has been selected for detailed study in this paper (Figure 3). This area exhibits a complicated spatial pattern of land cover and land use, providing a stringent test for both conventional and alternative classification techniques.

RESULTS AND DISCUSSION

Initial Land-Cover Classification

The first stage in SPARK processing is to segment the image into regions of uniform land cover. A variety of techniques can be used for this purpose, including unsupervised multispectral classification (Mather 1987), region growing (Rafat and Wong 1988, Li and Muller 1991) and split-and-merge techniques (Chen and Pavlidis 1979). A supervised, maximum-likelihood classification algorithm has been used here, because it is felt to offer the greatest control over both the number and nature of the classes defined. In this respect, six main land-cover types have been identified in the Orpington sub-scene, namely BUILT, TREE, CROP, GRASS, SOIL and WATER. The BUILT class correspond to roads and buildings within the main urban area; no attempt has been made to distinguish between these two surfaces in the present study. The remaining classes are largely self-explanatory. However, it is worth noting that the GRASS class incorporates regions of open space within the urban area (i.e. recreational land), as well as fields of permanent pasture outside it. Similarly, the CROP class incorporates, and is dominated by, areas of cultivated wheat and barley; again, no attempt has been made here to distinguish between these two crops.

After detailed visual examination of the digital image and 1:10,000 scale Ordnance Survey topographic maps of the corresponding area, a seventh land-cover type was identified. This class, which will be referred to as STRUCTURE, corresponds to large buildings such as factories, warehouses and hospitals. These exhibit a spectral response that is markedly different from other man-made structures in the remainder of the scene (houses and roads), presumably as a result of differences in the roofing materials used.

A series of irregularly-shaped regions, sampled systematically from within the image, have been used to define several training areas for each of the candidate land-cover classes. A similar set of regions have been used to define an independent test set. Some difficulty was experienced in creating suitable training and testing sets for the BUILT, WATER and STRUCTURE classes. In the case of the WATER and STRUCTURE classes this was because of their relatively limited areal extent; while for the BUILT class it was primarily due to the comparatively narrow, elongated regions that it forms. Consequently, the number of pixels used to train and test these classes is quite small.

A very low rejection threshold was set for the maximum-likelihood algorithm (>5 standard deviations, i.e. <0.001% pixels rejected) to produce an image with the minimum number of unclassified pixels. This is because an adjacency involving a NULL class pixel (referred to here as a NULL-adjacency) presents a problem at the spatial re-classification stage. Specifically, a NULL-adjacency may obscure the true spatial pattern of land-cover types present within the kernel. Thus, a BUILT-NULL adjacency may, in reality, represent BUILT-BUILT or BUILT-GRASS. Moreover, where there is more than one NULL-adjacency within the kernel, it may be impossible to determine the land use at that location. Thus, the errors introduced by using a very low rejection threshold must be balanced against the uncertainty introduced into the spatial re-classification procedure by a NULL-adjacency.

The results of the initial classification of land cover in the Orpington sub-scene are presented in Table 1 and Figure 4. Not surprisingly, given the limited number of broad land-cover classes used, a very high level of classification accuracy (Overall accuracy = 97.29%, Kappa coefficient = 0.93) has been achieved. The use of contiguous blocks of pixels for the test set probably mean that these values over-estimate the true level of accuracy achieved. However, visual comparison of Figures 3 and 4 suggests that this effect may not be too great.

Inferring Land-Use Information using SPARK

Having derived a satisfactory land-cover classification (Figure 4), SPARK has been used to re-classify this image into eight land-use categories; namely medium-density residential, low-density residential, commercial/industrial, woodland, arable farmland, permanent pasture, bare soil/fallow land and open water. The distinction made here between the medium-density and low-density residential categories is somewhat subjective. However, for the purpose of this study, medium-density housing broadly corresponds to terraced buildings with relatively small gardens, whereas low-density housing corresponds to detached and semi-detached buildings with somewhat larger gardens. Template vectors have been derived for each of the land-use categories using blocks of pixels selected at random from within larger sample areas. Several template vectors have been created in this way for each land-use category. These have been used to define a set of independent templates for that class. The accuracy of the re-classification has been tested using an independent set of irregularly-shaped sample areas. The land use in each of these areas has been determined from recent Ordnance Survey 1:10,000 scale maps, and verified through field checks.

The results obtained from SPARK using a 3x3 pixel kernel are given in Figure 5 and Table 2. These indicate that SPARK performs very well for the non-urban land-use categories, although this is to be expected given that they each comprise a single land-cover type. By contrast, the level of accuracy achieved for the low-density residential areas is very poor. Specifically, there is considerable confusion between this land

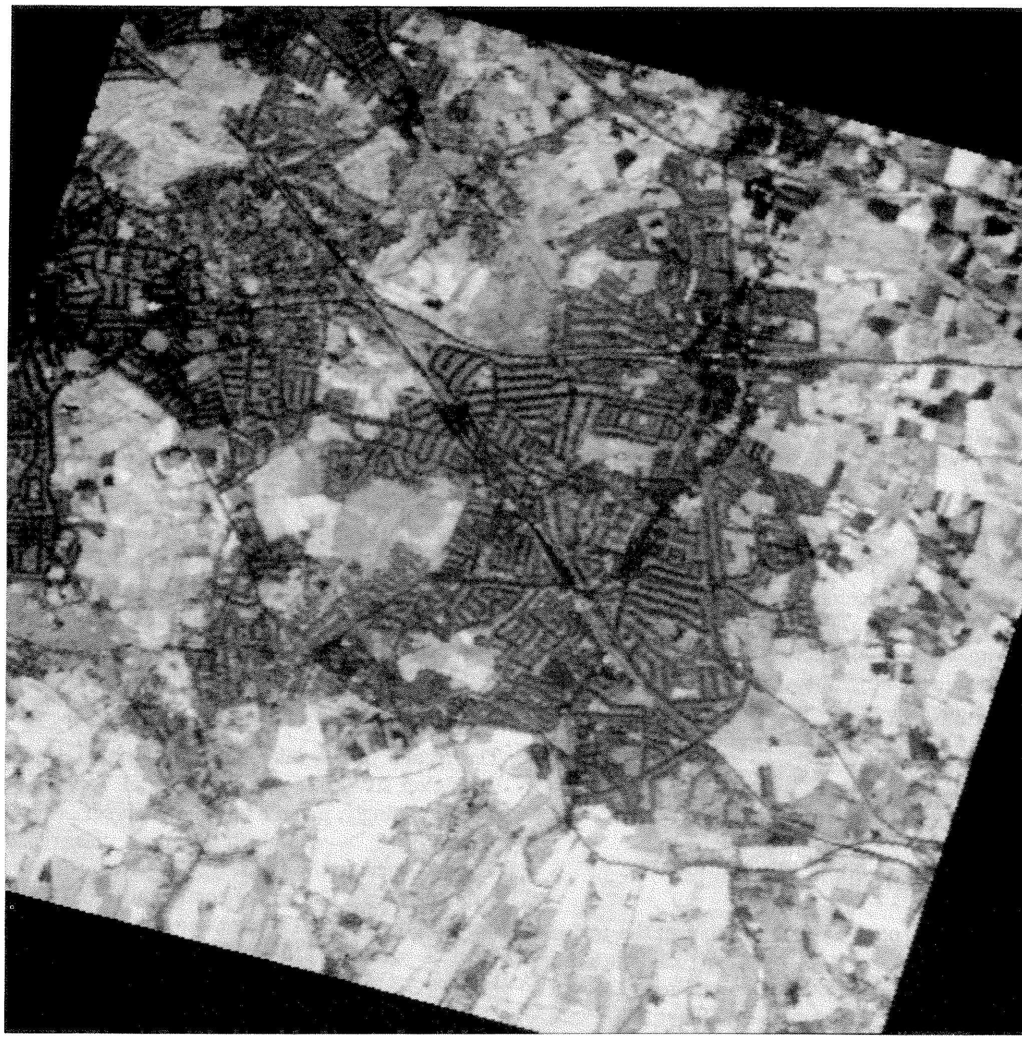

Figure 1 SPOT-HRV XS3 (Near-Infrared) image of Orpington in the London Borough of Bromley. © SPOT Image 1986, CNES.

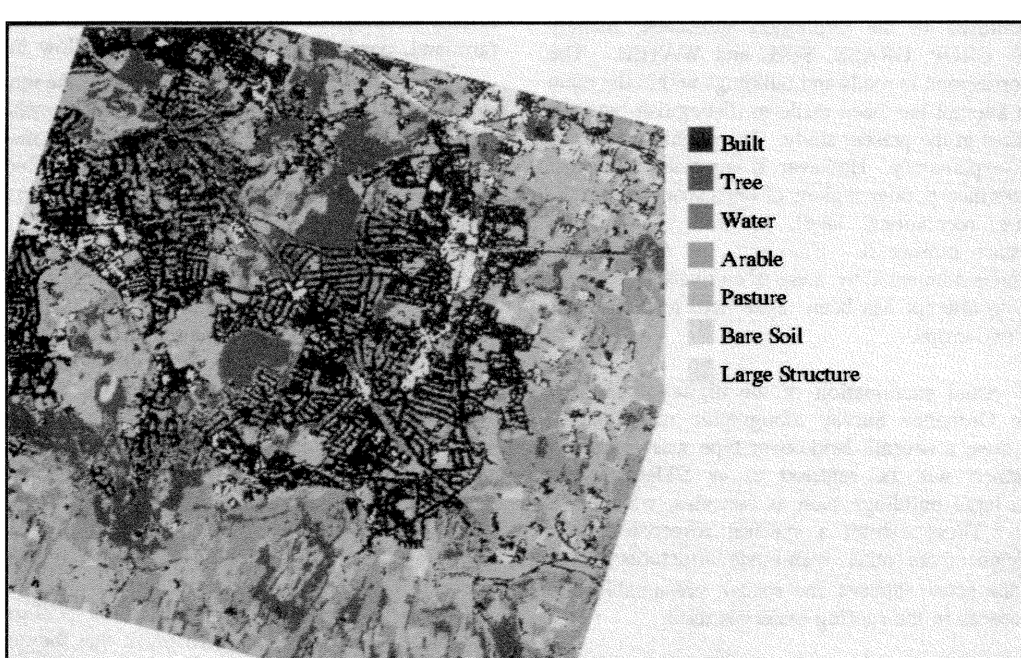

Figure 2 Land cover classification of SPOT-HRV multispectral image.

Table 1 Confusion Matrix for Per-Pixel Land Cover Classification

LAND COVER TYPE True ▸ Image ▾	Built	Structure	Tree	Crop	Grass	Soil	Water
Built	25	0	0	0	0	0	0
Structure	0	20	0	0	0	0	0
Tree	0	0	156	0	0	0	0
Crop	0	0	7	73	6	0	0
Grass	0	0	0	0	191	0	0
Soil	0	3	0	0	0	91	0
Water	0	0	0	0	0	0	18
Number of Test Pixels	25	23	163	73	197	91	18

Average Accuracy	97.09%	Overall Accuracy	97.29%	Kappa Coefficient (x100)	93.0%

Table 2 Confusion Matrix for Land Use Re-Classification Using SPARK (3 by 3 Pixel Kernel)

LAND USE True ▸ Image ▾	Low Density Residential	Medium Density Residential	Commercial/ Industrial	Woodland	Arable	Pasture	Fallow/ Bare Soil	Water
Low Density Residential	541	13	5	0	0	5	3	5
Med. Density Residential	695	625	1	0	0	0	0	0
Commercial/Industrial	0	0	17	0	0	0	0	0
Woodland	41	0	0	464	0	0	0	0
Arable	13	0	0	0	73	0	0	0
Pasture	0	2	0	0	0	192	0	0
Bare Soil	0	0	0	0	0	0	88	0
Water	0	0	0	0	0	0	0	7
No. of Test Pixels	1280	640	23	464	73	197	91	12

Average Accuracy	93.10%	Overall Accuracy	85.30%	Kappa Coefficient (x100)	43.96%

Table 3 Confusion Matrix for Land Use Re-Classification Using SPARK (9 by 9 Pixel Kernel)

LAND USE True ▸ Image ▾	Low Density Residential	Medium Density Residential	Commercial/ Industrial	Woodland	Arable	Pasture	Fallow/ Bare Soil	Water
Low Density Residential	1276	1	0	0	0	0	0	0
Med. Density Residential	4	609	0	0	0	0	0	0
Commercial/Industrial	0	0	23	0	0	0	0	0
Woodland	0	0	0	464	0	0	1	0
Arable	0	0	0	0	73	0	0	0
Pasture	0	0	0	0	0	197	0	0
Bare Soil	0	0	0	0	0	0	81	0
Water	0	0	0	0	0	0	9	12
No. of Test Pixels	1280	640	23	464	73	197	91	12

Average Accuracy	97.98%	Overall Accuracy	96.86%	Kappa Coefficient (x100)	92.1%

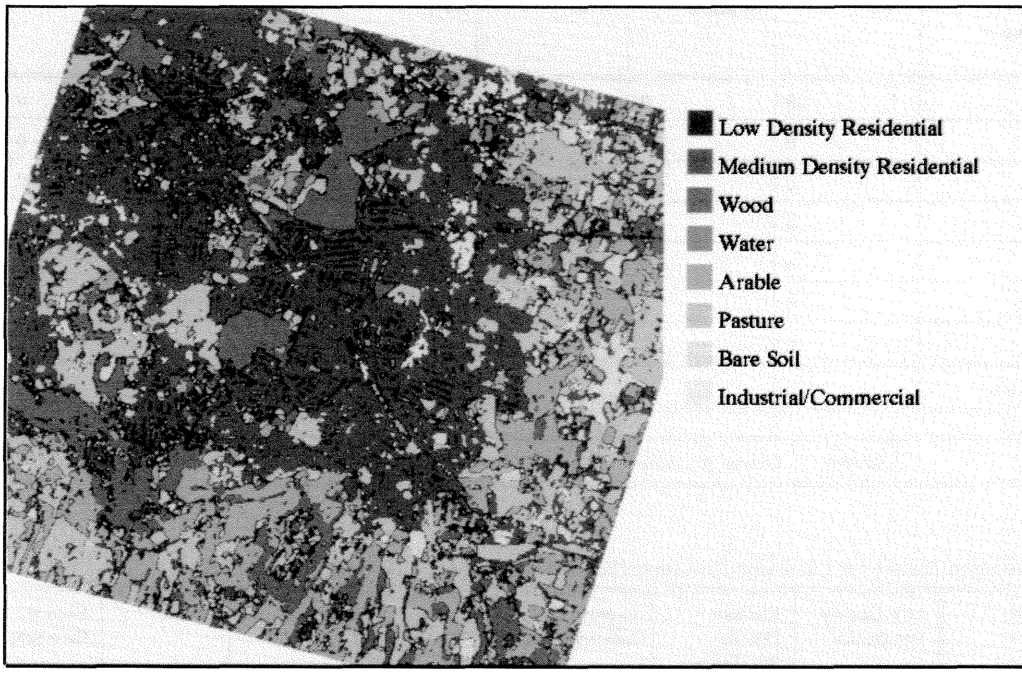

Figure 3 Land-use classification using SPARK (3x3 pixel kernel).

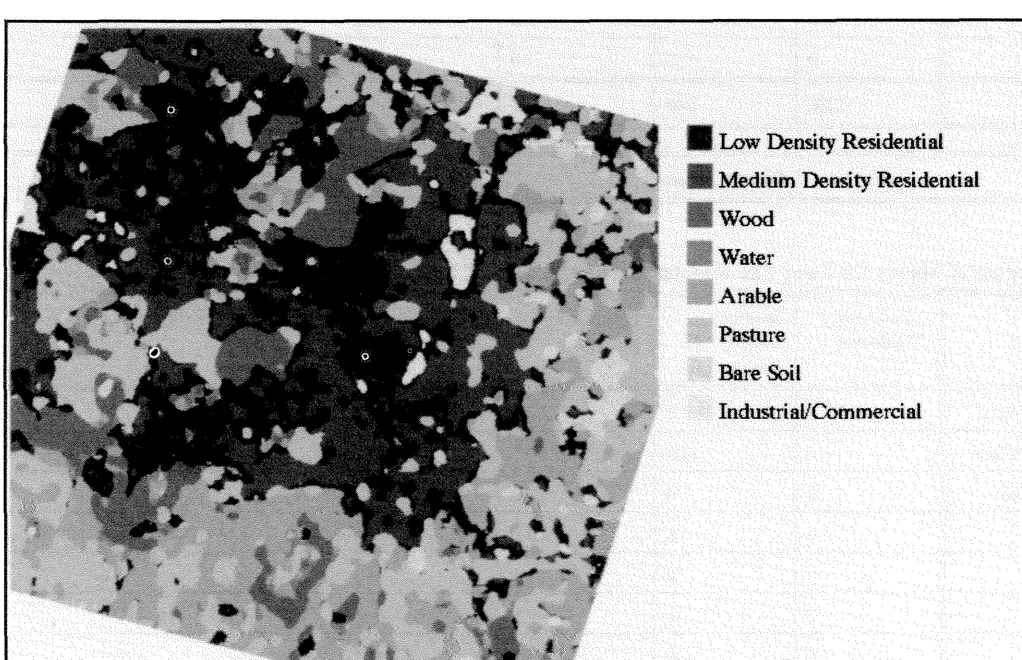

Figure 4 Land-use classification using SPARK (9x9 pixel kernel).

use and the medium-density residential areas. A close examination of Figures 3 and 5 reveals that this occurs where individual streets are very widely spaced, such that the roads and some of the buildings (i.e. the BUILT land-cover) are assigned to the medium-density residential category, while buildings which adjoin some form of open space (e.g. large gardens) are assigned to the low-density residential category. Consequently, the street pattern is still apparent in some low-density residential districts. These results suggest that a 3x3 pixel kernel is too small to take into account the spatial pattern of land-cover types typical of some urban land-use categories.

In view of the results obtained above, tests have been performed on SPARK using a range of other kernel sizes. Only the results obtained using SPARK with a 9x9 pixel kernel will be presented here (Figure 6 and Table 3). These indicate a substantial improvement in the accuracy with which the candidate land-use categories have been classified (Overall accuracy = 96.86%, Kappa coefficient = 0.9209). This is chiefly a result of a reduction in the confusion between the low- and medium-density residential zones. Nevertheless, some problems are apparent in Figure 6. In particular, the boundaries between land parcels outside the urban area appear to have been smoothed. Since this did not occur using the 3x3 pixel kernel, it suggests that different kernel sizes may be appropriate for use in different parts of the image.

ENHANCEMENTS AND FURTHER STUDY

Clearly, any technique which attempts to infer land use by examining the spatial pattern of land cover in an image will be sensitive to the accuracy of the initial land-cover classification. Apart from examining alternative methods of image segmentation (e.g. region-growing and split-and-merge algorithms), further study is required to determine the sensitivity of the spatial re-classification to the number and nature of the candidate classes identified at this stage.

The basic SPARK procedure could also be improved by including information on the 'likelihood' with which each pixel is assigned to a particular land-cover class. At present, the label associated with each pixel is implicitly assumed to be correct and without error. In reality, we may only have a low confidence in that label. Consequently, there must also be some uncertainty about the type of adjacency (i.e. Built-Built, Built-Structure,...) that exists between neighbouring pixels. This can be taken into account by calculating the product of the 'likelihood' values associated with their class labels. This would give a probability value for each adjacency which, when summed for all pairs of pixels within the kernel, could be used when comparing the measured adjacency vector with those of the land-use templates.

SUMMARY AND CONCLUSIONS

A technique has been developed that attempts to infer land use in urban areas from satellite sensor images previously classified into discrete land-cover types. The procedure examines the spatial arrangement of class labels within a moving window and compares this with that found in areas of the image for which the land use is known. Preliminary tests of this approach indicates that very high levels of accuracy can be achieved, both in terms of the initial land-cover classification (>95%) and the subsequent re-classification of land-use (>85%). A series of further enhancements to the method have been suggested.

Results obtained from this technique appear to be sensitive to the size of kernel employed. Different sizes of kernel have been found to be appropriate in separate parts of the scene: while a large kernel is required to represent the typical spatial arrangement of land-cover types in urban areas, application of the same kernel to agricultural areas can produce blurring at field boundaries. Although it may be possible to develop procedures to vary the size of kernel applied to different parts of the image, this is likely to introduce substantial computational overheads.

A more serious criticism is that kernel-based re-classification imposes arbitrary and unhelpful limits on the area over which the spatial mixing of land-cover types is investigated. Consequently, alternative techniques for spatial re-classification are also being investigated (Barr 1992). These operate by identifying specific land-cover 'objects' (i.e. regions of uniform land cover). Land use is inferred by examining the spatial relationships between these objects.

ACKNOWLEDGEMENTS

The authors would like to thank both the Economic and Social Research Council and the Natural Environment Research Council for support through a research grant under the joint ESRC/NERC Geographic Data Handling programme. The authors are particularly grateful to Graham Sadler and Andrew Williams for detailed discussions during the early stages of this project. Thanks are also due to Asad Hamid, Prof. Peter Muller, Dr. Allison Reid, (UCL) and Prof. Paul Mather (Nottingham) for their comments and ideas.

REFERENCES

Atkinson, P., Cushnie, J.L., Townshend, J.R.G., and Wilson, A.K., 1985, Improving Thematic Mapper land cover classification using filtered data, *Int. J. Rem. Sens.*, **6**, 955-961.

Barnsley, M.J., Sadler, G.J., and J.S. Shepherd, 1989, Integrating remotely-sensed images and digital map data in the context of urban planning, *Proc. 15th Ann. Conf. Remote Sensing Society*, 13-15th September 1989, Bristol, U.K., (Nottingham: Remote Sensing Society), 25-32.

Barnsley, M.J., Barr, S.L., Sadler, G.J, 1991. Spatial Re-Classification of Remotely sensed Images for Urban Land Use Monitoring. *Proc. Spatial Data 2000*, 17-20 September 1991, Oxford University, (Nottingham : Remote Sensing Society), 106-117.

Barr, S.L., 1992, Object-based re-classification of high resolution digital imagery for urban land-use monitoring. These Proceedings.

Chen, P.C., and Pavlidis, T., 1979, Segmentation by texture using a co-occurrence matrix and a split-and-merge algorithm, *Computer Vision, Graphics and Image Processing*, **10**, 172-182.

Franklin, S.E., and Peddle, D.R., 1990, Classification of SPOT-HRV imagery and texture features, *Int. J. Rem. Sens.*, **11**, 551-556.

Forster, B.C., 1980, Urban residential ground cover using Landsat digital data, *Photogramm. Engng. & Rem. Sens.*, **46**, 547-558.

Forster, B.C., 1985, An examination of some problems and solutions in monitoring urban areas from satellite platforms, *Int. J. Rem. Sens.*, **6**, 139-151.

Gastellu-Etchegorry, J.P., 1990, An assessment of SPOT XS and Landsat MSS data for digital classification of near-urban land cover, *Int. J. Rem. Sens.*, **11**, 225-235.

Gong, P., and Howarth, P.J., 1989, Performance analyses of probabilistic relaxation methods for land-cover classification, *Rem. Sens. Environ.*, **30**, 33-42.

Gong, P., and Howarth, P.J., 1990, The use of structural information for improving land-cover classification accuracies at the rural-urban fringe, *Photogramm. Engng. & Rem. Sens.*, **56**, 67-73.

Gong, P., and Howarth, P.J., 1992, Frequency-based contextual classification and gray-level vector reduction for land-use identification, *Photogramm. Engng. & Rem. Sens.*, **58**, 423-437.

Guo, Liu Jian, and McM. Moore, J., 1991, Post-classification processing for thematic mapping based on remotely-sensed image data, In *Proc. Int. Conf. IEEE Geoscience and Remote Sensing Society*, 3-7 June 1991, Espoo, Finland, (New York: IEEE), 2203-2206.

Gurney, C.M., 1981, The use of contextual information to improve land cover classification of digital remotely sensed data, *Int. J. Rem. Sens.*, **2**, 379-388.

Gurney, C.M. and Townshend, J.R.G., 1983, The use of contextual information in the classification of remotely sensed data, *Photogramm. Engng. & Rem. Sens.*, **49**, 55-64.

Haack, B., Bryant, N., and Adams, S., 1987, An assessment of Landsat MSS and TM data for urban and near-urban land-cover digital classification, *Rem. Sens. Environ.*, **21**, 201-213.

Haralick, R.M., 1979, Statistical and structural approaches to texture, *Proceedings of the IEEE*, **67**, 786-804.

Jackson, M.J., Carter, P., Smith, T.F., and Gardner, W., 1980, Urban land mapping from remotely sensed data, *Photogramm. Engng. & Rem. Sens.*, **46**, 1041-1050.

Li, K., and Muller, J-P., 1991, Segmenting satellite imagery: A region growing scheme, *Proc. Conf. IEEE Geoscience and Remote Sensing Society (IGARSS'91)*, Finland 1991.

Martin, L.R.G., Howarth, P.J., and Holder, G., 1988, Multispectral classification of land use at the rural -urban fringe using SPOT data, *Canadian J. Rem. Sens.*, **14**, 72-79.

Mather, P.M., 1987, *Computer Processing of Remotely-Sensed Images*, (London:John Wiley), 289-309.

Rafat, H.M., and Wong, A.K.C., 1988, A texture information-directed region growing algorithm for image segmentation and region classification, *Computer Vision, Graphics and Image Processing*, **43**, 1-21.

Sadler, G.J., and Barnsley, M.J., 1990, Use of population density data to improve classification accuracies in remotely-sensed images of urban areas, *First European GIS Conference (EGIS'90)*, Amsterdam, The Netherlands, 10-13 April 1990, (Utrecht:EGIS Foundation), 968-977.

Sadler, G.J., Barnsley, M.J., and Barr, S.L., 1991, Information extraction from remotely-sensed images for urban land analysis, *Second European GIS Conference (EGIS'91)*, Brussels, April 1991, (Utrecht:EGIS Foundation), 955-964.

Wharton, S.W., 1982, A contextual classification method for recognizing land use patterns in high resolution remotely-sensed data, *Pattern Recognition*, **15**, 317-324.

Whitehouse, S., 1990, A spatial land-use classification of an urban environment using high-resolution multispectral satellite data, In *Proc. Conf. on Remote Sensing and Global Change*, University College Swansea, 19-21 September 1990, (Nottingham: Remote Sensing Society), 433-437.