Segment-Based Land-Use Classification from SPOT Satellite Data

Katarina Johnsson

Abstract
Spectrally heterogeneous land-use categories cannot be adequately classified from high resolution satellite data, using conventional multispectral classification techniques. In this study, built-up land is extracted based on the spectral and the spatial properties of the segments in a spectrally classified satellite image. Object structures and classification rules have been implemented in an expert system shell – Nexpert Object. The approach has been tested using SPOT multispectral data covering an area south of Stockholm, Sweden. The classification accuracy for built-up land improved significantly after the application of a few relatively simple rules, based on segment size and on relations between segments. The object oriented implementation in Nexpert Object is flexible but slow, and performance problems were encountered due to the large number of objects being processed.

Introduction
Land use is one of the basic data layers in geographic information systems (GIS) for physical planning and environmental modeling. High resolution satellite data are often mentioned as an important source of input, e.g., to update existing digital land-use maps. Traditional multispectral image classification techniques are, however, insufficient for extraction of complex and spectrally heterogeneous land-use categories, such as built-up land, from high resolution satellite data [e.g., Baker *et al.*, 1991; Sadler, 1991].

Spatial information can be incorporated to improve the classification accuracy for spectrally heterogeneous categories [e.g., Gonzalez and Wintz, 1987; Lillesand and Kiefer, 1987]. One way to introduce spatial information is by segmenting an image and computing spatial attributes for each segment. Typical attributes may be area, length of perimeter, compactness (area/perimeter^2), degree and type of texture, or minimum bounding rectangle [Patterson, 1990; Gonzalez and Wintz, 1987]. This method of describing the spatial properties of image objects was originally developed within the field of computer vision [e.g., Gonzalez and Wintz, 1987], where it has been used as a component in knowledge-based image analysis systems, i.e., systems developed to identify and describe objects and object relationships in images [Patterson, 1990].

Within the field of remote sensing, knowledge-based image analysis has mainly been applied to the analysis of aerial images. The work most often cited in this context is by Nagao and Matsuyama [1980]. Nagao and Matsuyama developed an expert system for classification of multispectral aerial images, based on characteristic regions. The regions were defined by their spectral and spatial features. Production rules were used to implement the classifier. Another example of rule-based land-use classification of an aerial photograph, based on the spectral and spatial features of image segments, is presented by Mehldau and Schowengerdt [1990]. In addition to using some of the spatial descriptors listed above, Mehldau and Schowengerdt incorporated image context in the rule-based classification. The properties of neighboring segments were evaluated and used to increase or decrease the confidence values of the classified segments. Mehldau and Schowengerdt comment that this was the potentially most powerful phase of the classification process.

Goodenough *et al.* [1987] used a knowledge-based approach, incorporating spatial segment features, to map-image congruency evaluation, using Landsat TM data. Civco (1989) developed an expert system for land-use and land-cover mapping, based on Landsat TM data and an existing land-use/land-cover map. Civco uses an existing land-use map to delineate image regions with the same land use/land cover. He assumes that these regions are spectrally homogeneous and extracts a number of spectral and spatial attributes for each region. The regions and their attributes are input to a rule-based expert system developed in the EXSYS expert system shell. Civco achieves a classification accuracy of 100 percent, partly because he develops case-specific rules to correct the misclassifications that occur with the original set of rules. The segmentation methodology also had an influence on the accuracy.

Argialas and Harlow [1990] give an extensive overview over various approaches to image classification and image understanding within remote sensing. Their article also contains a review of knowledge-based image analysis systems developed for remote sensing.

Aim of study
The aim of this study has been to implement and apply a knowledge-based approach to classification of built-up land from multispectral SPOT data. Previous work, as described above, has served as a source of ideas and inspiration. This study differs from the ones described above in that it addresses the specific problem of extraction of a spectrally heterogeneous land-use category from SPOT XS satellite data. The study is based on the assumption that built-up land, being a complex and spectrally heterogeneous land-use category, can be described, not by its typical overall spectral response, but by its composition of small segments corresponding to a number of different land-cover classes. The study also addresses the advantages and problems related to the imple-

Photogrammetric Engineering & Remote Sensing,
Vol. 60, No. 1, January 1994, pp. 47–53.

0099-1112/94/6001–47$03.00/0

Department of Photogrammetry, Royal Institute of Technology, S-10044 Stockholm, Sweden

mentation of data structures and classification rules in an object-oriented, rule-based expert system shell – Nexpert Object [Neuron Data, 1990].

Test Area and Data

The test area is located approximately 20 km southwest of Stockholm, Sweden (Figure 1). It is an area of 10 by 10 km² found on the 1:50.000-scale topographic map sheet Stockholm 10 I SV. The area corresponds to a 512- by 512-pixel SPOT multispectral image. The SPOT scene was registered on 8 May 1986 (K/J 061-229 860508). The image has been radiometrically and geometrically corrected (to the Swedish National Grid).

The test area is close to the city of Stockholm and is, therefore, subject to intensive development. Approximately 25 percent of the area is built-up land, most of it residential. Two major types can be distinguished – areas with apartment buildings and areas with single family houses. There are also some industrial areas. The relief within the test area is less than 100 m. The vegetation was not yet fully developed at the time of the satellite data registration .

The data set also includes a digital land-use map and an orthophoto, both used for comparison and evaluation of the classification accuracy. The land-use map was created around 1975 by aerial photointerpretation (infrared photographs). The digitized map sheet is stored in raster format. The exact method and accuracy of the digitizing procedure is not known. The map contains 63 land-use categories and the original pixel size is 50 by 10 m. The orthophoto (scale 1:50,000) is from early May 1986, i.e., it was registered approximately at the same time as the satellite scene.

The study has been carried out at the Department of Photogrammetry, KTH, Stockholm, on a UNIX based HP Apollo workstation with 8 Mbyte of random access memory.

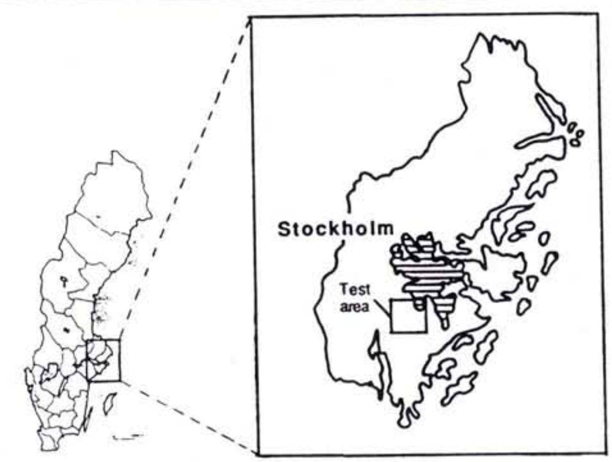

Figure 1. The test area is located approximately 20 km southwest of Stockholm, Sweden.

Methods

The classification procedure involves spectral classification, determination of size and neighbor relations for the segments in the spectrally classified image, and rule-based classification of the image segments into land-use categories (Figure 2). In the rule-based classification phase, the image segments are treated as objects with a set of properties. Thus, the procedure also involves a conversion from raster format to a list of objects and vice versa. The rule based classifier has been implemented in the Nexpert Object expert system shell, as well as in a conventional programming language. The performance of the two implementations were compared by measuring the processing time for sub-images over the test area. Accuracy evaluation was carried out using the existing land-use map as "ground truth." The effect of the classification rules was estimated by comparing the classification accuracy in the final map with the classification accuracy of the spectrally classified image.

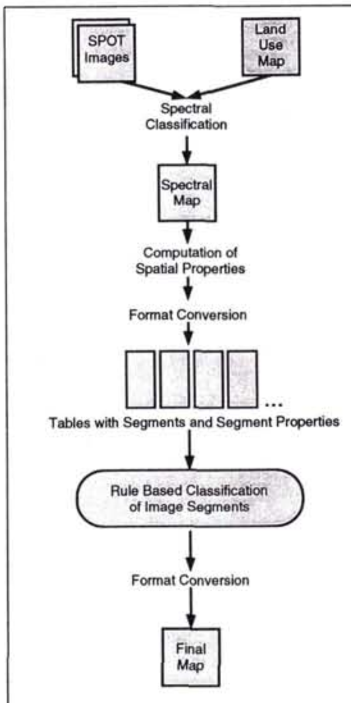

Figure 2. Segment-based classification of spectrally heterogeneous land-use categories.

Spectral Classification and Computation of Spatial Segment Descriptors

The initial spectral classification of the image was performed by a conventional maximum-likelihood classification, using the red and infrared bands of the SPOT image (bands 2 and 3). The classification was guided by the existing land-use map, i.e., the existing map was used as a basis for a semi-automatic selection of training data. For further details on the methodology for training area selection and image classification see Johnsson [1990]. Due to software limitations, only two image bands could be used. It was shown that band 1 could be excluded without serious loss of information [Johnsson, 1991] and other authors have described a high correlation between band 1 and band 2 [e.g., Sadler, 1991]. The spectral classification resulted in an image with eight spectral classes, labeled water, coniferous forest, deciduous (and mixed) forest, grass (and other herbaceous plants), bare soil, urban 1, urban 2, and urban 3 (Figure 3).

Size, perimeter, and a list of neighbors were computed for each segment, using programs implemented in C and Pascal [Huang, 1991; Johnsson and Kanonier, 1991]. These spatial parameters, along with the spectral class and a unique identity number for each segment, were stored in tables. The tables were implemented as ASCII files. One set of tables contains the spectral class, size, and perimeter of each segment and another set of tables describes the relations between

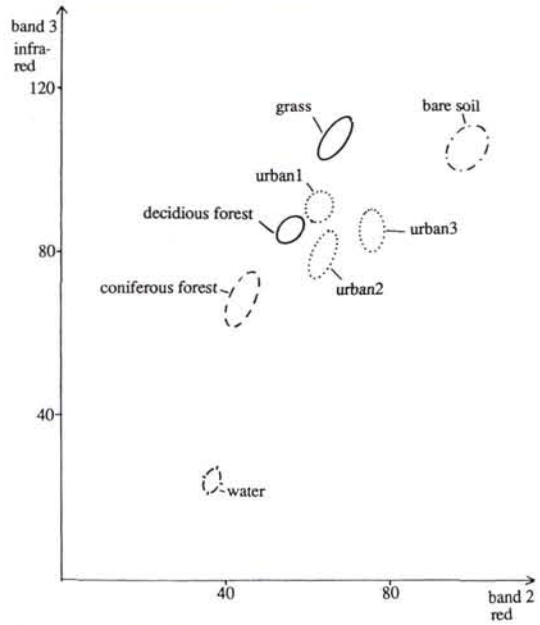

Figure 3. Spectral classes after a supervised maximum-likelihood classification based on the red and the infrared band of the SPOT satellite data.

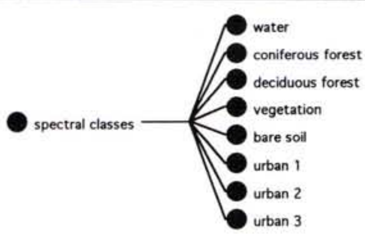

Figure 4. The class structure representing spectral classes.

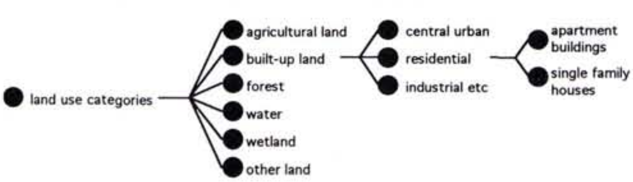

Figure 5. The class structure representing land-use categories.

neighboring segments. Thus, the available spectral and spatial information for the segments has been collected in tables on which a rule-based classifier can operate.

Classification of Image Segments in the Nexpert Object Expert System Shell

Nexpert Object is a hybrid expert system shell, which supports an object-oriented (semantic net) representation dimension, as well as a rule-based reasoning dimension [Neuron Data, 1990]. Mechanisms for inheritance between classes and objects are supplied. Procedures may be incorporated in the objects through so called *meta-slots*, attached to the object properties. The inference engine is agenda based. *Dynamic objects* can be created at runtime, e.g., by a rule, and they are deleted from the knowledge base at the end of a session. The rules in Nexpert Object are symmetric, i.e., they can be evaluated in both the forward and backward direction. *Backward chaining* of rules by default has the highest priority in the agenda. The rules work on the objects, or, in the case of generalized rules, on the classes. (For further details of knowledge representation in expert systems, see, for example, Harmon and King (1985), Genesereth and Nilsson (1988), or Winston (1984).)

In this study, the following features of the Nexpert Object Expert system shell were used:

- The *object oriented data representation scheme*, to represent the original spectral classes (Figure 4) and the hierarchy of land-use categories (Figure 5).
- The *dynamic objects concept*, to represent the actual image segments and their properties (Figure 6), as extracted from the ASCII tables. Dynamic objects were also used to represent relations between neighboring segments (Figure 7).
- The *ability to link dynamic objects to existing object structures*, as a means of classifying the dynamic objects (i.e., the image segments). Initially, a dynamic object is created for each segment that is input for classification. The object is also linked to its corresponding spectral class in the object structure. As rules based on spatial properties, i.e., size and relations to neighboring objects, are applied, classification is carried out by linking the dynamic objects to the appropriate class in the land-use category structure. Nexpert Object supports multiple class assignment. This makes it possible for an object to remain linked to the spectral class structure, even after it has been linked to another parent class in the land-use category structure. It would also be possible to link an object to several different potential land-use categories, as an intermediate step, before the final evaluation is made as to which land-use category is the most likely one. However, this study never progressed as far as to make use of this possibility.
- The *production rules*, to implement the processing strategies and the classification criteria. The rules were mainly evaluated through backward chaining.
- The *object meta slots*, to implement the classification of an object (segment) based on the properties of the neighboring segments.

All image segments are input to the expert system, when the classifier is invoked (Figure 8). The segments and their properties are represented as dynamic objects, on which the rules can operate. The neighbor relations are input only if

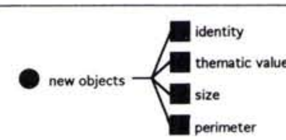

Figure 6. The class new objects and its properties.

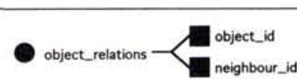

Figure 7. The class structure representing relations between objects.

and when a rule requires them (thus, the implementation as a meta-slot).

Accuracy Evaluation

The simplest solution for accuracy evaluation was to use the existing land-use map as a control data layer. For this purpose it was simplified (reclassified) into seven categories (Table 1). It was also resampled to the same pixel size as the satellite data (20 by 20 m).

The following properties of the existing land-use map should be kept in mind:

- It is not up-to-date;
- The accuracy of the original data is not well known;
- There has been a loss of information and a decrease in geometric accuracy during the processing of the data.

The classification accuracy was evaluated by a pixel-wise comparison of the classified image to the existing digital land-use map in a raster-based geographic information system (GRASS) [Westervelt, 1988]. The coincidence between the categories in the two data layers was computed, using all pixels in the maps. The coincidence matrices were summarized using three accuracy measures . *Class accuracy* is the number of correctly classified pixels divided by the number of pixels belonging to that particular class in the classified map. *Category accuracy* is the number of correctly classified pixels divided by the number of pixels belonging to that particular category in the control data layer. *Mean accuracy* is the mean of the two other measures.

When the classified images were compared to the existing land-use map, the following combinations were used (the first element in a pair is the image spectral class and the second is the land-use category): *water/water, coniferous forest/coniferous forest, deciduous forest/deciduous forest, bare soil/agricultural land, grass/grassland,* and *urban (all urban classes merged)/built-up land*. The comparison between bare soil and agricultural land was possible as most of the fields were barren at the time of the satellite registration.

Built-up land was also digitized from the orthophoto. The photo is up-to-date in relation to the satellite data. The spatial and spectral properties of the orthophoto did not allow a detailed interpretation. Therefore, the areas of built-up land were only roughly delineated. The digitized boundaries were stored in vector format, imported to GRASS, and overlaid on the classified images for visual comparison.

Areas that have been exploited between the creation of the existing land-use map and the satellite registration were detected by displaying the new vector data together with the existing map. These areas were stored as a separate information layer in the GRASS raster database and area statistics were computed. The area of built-up land has increased approximately 10 percent between 1975 and 1986.

From a comparison between built-up land as digitized from the orthophoto and built-up land in the existing land-use map, it was estimated that the uncertainty in the area coverage of built-up land in the old land-use map may be as high as 20 percent [Johnsson, 1991], reasons being errors in the original photointerpretation, digitizing errors, errors introduced in the pre-processing of the land-use map, and possibly other errors as well.

Experimental Classification

Classification Rules

After multispectral classification, the image consisted of eight spectral classes, three of them related to man-made objects (urban1-3). All spectral classes, except water, existed within built-up areas, mainly as small segments, corresponding to real world objects such as groups of trees (forest), gardens (grass), and parking lots (bare soil). The existing land-use map was used as a template to extract those areas in the classified image map that correspond to built-up land. The size distribution of segments belonging to the different spectral classes was determined and used as a basis for the thresholds in the rules. Four experimental classification rules were designed and implemented in the expert system as well as in a conventional Pascal module.

A segment is part of *built-up land*

(1) IF the segment belongs to the spectral class *bare soil* and is smaller than 100 pixels, OR;
(2) IF the segment belongs to the spectral class *coniferous forest* and is smaller than 100 pixels, OR;
(3) IF the segment belongs to the spectral class *deciduous forest* and at least one of the neighboring segments belongs to any of the *urban* spectral classes, OR;
(4) IF the segment belongs to the spectral class *grass* and is smaller than 200 pixels.

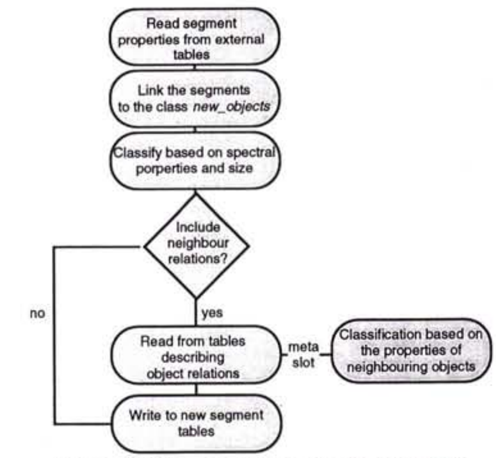

Figure 8. Process flow in Nexpert Object.

Table 1. Categories in the Simplified Land-Use Map for Accuracy Evaluation.

1.	Water
2.	Coniferous forest
3.	Deciduous and mixed forest
4.	Agricultural land
5.	Grassland
6.	Built-up land
7.	Other land use

Classification Results

Figure 9 shows how the four rules affect a built-up area, Tumba. Figure 9a is the original spectral classification; Figures 9b-9e are the results after applying the four rules, respectively; Figure 9f is the combined effect of the four rules; and, finally, Figure 9g is the existing land-use map for comparison. From these images it can be concluded that the classification rules indeed do improve the classification of built-up land. Built-up land in Figure 9f is an area-covering, continuous category, whereas in Figure 9a it is a heterogeneous mixture of various spectral classes.

Plate 1 shows the whole test area after application of the

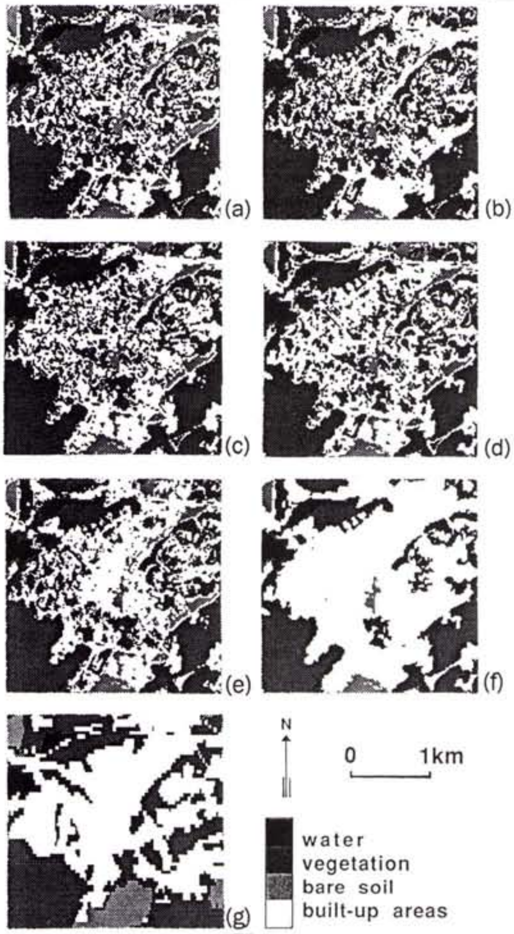

Figure 9. Tumba(a) Original spectral classification, (b) segments of bare soil smaller than 100 pixels have been classified into built-up land, (c) segments of coniferous forest smaller than 100 pixels have been classified into built-up land, (d) segments of deciduous forest with urban neighbors (urban 1, urban 2, or urban 3) have been classified into built-up land, (e) segments of grass smaller than 200 pixels have been classified into built-up land (f) a combination of b-e, (g) the existing land-use map.

four classification rules[1]. The white lines show *built-up areas* as delineated from the orthophoto. The pink color represents areas classified into built-up land by the rules (new urban). As can be seen, a large proportion of these areas actually fall within built-up land as delineated from the orthophoto. However, it is also obvious that the rules have a unwanted side effects. *Built-up land* is "created" where there is not supposed to be any.

The mean classification accuracy for *built-up land* has increased by 22 percent (from 45 percent to 55 percent) after application of the four classification rules (Table 2). This is mainly due to an increase in category accuracy by 80 percent (from 42 percent to 70 percent). This supports the conclusions, that the classification accuracy mainly has improved *within* the areas of *built-up land*. In fact, the class accuracy for *built-up land* has decreased after application of the four rules.

The mean accuracy for *agriculture/bare soil* has increased, indicating that the rule for segments of *bare soil* indeed makes a fairly good separation of segments of bare soil into the land-use categories *built-up land* and *agriculture*. The accuracy for *coniferous forest* remained more or less the same. The classification accuracy for *grassland/grass* and *agriculture/bare soil*, however, shows severe deterioration.

All four rules have a positive effect on the classification accuracy for *built-up land*, especially on the category accuracy (Table 3). However, three of them also cause a decrease in class accuracy. It is especially the case for the rule based on *deciduous forest* with *urban* neighbors.

It can be argued that the classification accuracy in gen-

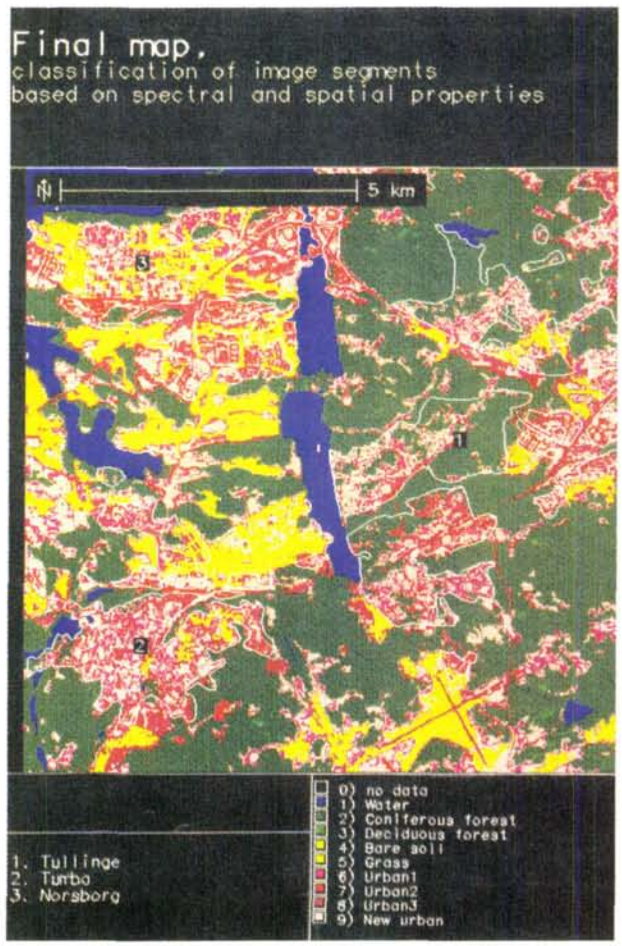

Plate 1. Final land-use classification after application of four segment-based (segment-based) classification rules.

[1] It was found that the number of segments in the test image exceeds the number of dynamic objects that can be created in Nexpert Object. Thus, only sub-images could be processed in the Nexpert Object classifier (see below). Classification of sub-images was carried out using both implementations and the results were identical. For reasons of efficiency, the Pascal module and not the Nexpert Object module was used to create the map covering the entire test area (Plate 1).

TABLE 2. ACCURACY EVALUATION BASED ON COINCIDENCE COMPUTATIONS BETWEEN THE EXISTING LAND-USE MAP AND THE CLASSIFIED MAP BEFORE AND AFTER THE APPLICATION OF THE FOUR CLASSIFICATION RULES.

Class	Class Accuracy (%) Before	Class Accuracy (%) After	Class Accuracy (%) Change	Category Accuracy (%) Before	Category Accuracy (%) After	Category Accuracy (%) Change	Mean Accuracy (%) Before	Mean Accuracy (%) After	Mean Accuracy (%) Change
water	97	97	0	92	92	0	94	94	0
coniferous forest	57	59	+4	82	79	−4	67	68	+1
deciduous forest	14	13	−3	12	1	−100	13	1	−100
agriculture/bare soil	46	68	+48	45	41	−9	46	51	+11
grassland/grass	18	16	−11	38	18	−53	24	17	−29
built-up land	50	45	−10	40	72	+80	45	55	+22

TABLE 3. INFLUENCE OF THE FOUR RULES ON THE ACCURACY OF THE LAND-USE CATEGORY BUILT-UP LAND. BASED ON COINCIDENCE COMPUTATIONS BETWEEN THE EXISTING LAND-USE MAP AND THE CLASSIFIED IMAGES.

Rule number	Class Accuracy (%) Built-Up Land Before	Class Accuracy (%) Built-Up Land After	Class Accuracy (%) Built-Up Land Change	Category Accuracy (%) Built-Up Land Before	Category Accuracy (%) Built-Up Land After	Category Accuracy (%) Built-Up Land Change	Mean Accuracy (%) Built-Up Land Before	Mean Accuracy (%) Built-Up Land After	Mean Accuracy (%) Built-Up Land Change
(1)	50	50	0	40	46	+15	44	48	+9
(2)	50	48	−4	40	45	+13	44	46	+5
(3)	50	45	−10	40	51	+28	44	48	+9
(4)	50	49	−2	40	51	+28	44	50	+14

eral is unusually low. This can to some extent be explained by the fact that nearly all pixels in the land-use map were used for accuracy evaluation, instead of just a sample. Due to geometric inaccuracies in the land-use map, there are zones of uncertainty and error along the boundaries between regions of different land-use categories. These zones are included in the accuracy evaluation, and a certain percentage of the misclassifications reflects these geometric inaccuracies more than actual misclassifications in the satellite data.

Performance of the Classifier Implemented in Nexpert Object

The spectrally classified image contains approximately 21,000 image segments. It turned out that Nexpert Object cannot handle that many dynamic objects. Empirical tests showed that the maximum possible number of dynamic objects is less than 11,000. Thus, only small images could be classified with this system. Quadrants of the test image (256 by 256 pixels) could be used when rules based on segment size were applied. When rules based on neighbor relations were applied as well, only small sub-images (50^2 to 100^2 pixels) could be processed.

Quadrants of the image were used as input to the knowledge base for classification based on spectral class and size. The quadrants (256 by 256 pixels) contain 5000 to 6000 segments each. It took 35 to 45 minutes to classify each quadrant (applying one rule based on spectral class and size) on an HP Apollo workstation with 8 Mbyte RAM, and no other processes running. Approximately 80 percent of that time was used for reading and writing from and to the external files containing the data tables. It took approximately 30 seconds to do the same thing, using the Pascal module.

Classification based on neighbors caused a dramatic increase in the number of dynamic objects. The third quadrant contains 5908 segments and the segments of one spectral class (deciduous forest) have 5089 relations, i.e., approximately 11,000 dynamic objects were needed. This could not be handled by the system. The neighbor based classification was tested on a 50- by 50-pixel sub-image over Tumba, one of the residential areas. This sub-image resulted in 553 segments and 1130 relations, i.e., approximately 1600 dynamic objects. It took approximately 30 minutes to classify the sub-image in Nexpert Object. The Pascal module required one to two minutes to do the same thing.

Discussion

Compared to the initial multispectral classification, the segment-based classification did improve the classification results. The main source of error is pixels that were erroneously classified as urban in the initial spectral classification. The errors were propagated and augmented by the rule based on urban neighbors. Segment size is a useful classification criterion for built-up land, as these areas mainly contain small segments. However, the rules based on segments size also introduced new errors, as not all small segments are part of built-up land. Size should be combined with other classification criteria. Systematic studies are needed to find and implement other useful segment descriptors.

The absolute values for classification accuracy may seem very low. To some extent this is an effect of the method and the material used for the accuracy evaluation. It is debatable to use a ten year old data layer for accuracy evaluation . However, the main purpose in this study was to get a relative measure of the classification accuracy, to be able to compare the results before and after the segment classification. The improvement in category accuracy for *built-up land* from 40 to 72 percent is large enough to be significant, even considering the uncertainties in the existing land-use map, and so is the increase in class accuracy for the combination *agriculture/bare soil*. For the other classes, the change in classification accuracy could not be proved to be significant. In these cases, the effect of the classification rules could mainly be judged from the visual comparison between the digitized boundaries of built-up land in the orthophoto.

The classification rules were implemented both in the Nexpert Object expert system shell and as Pascal modules. The Nexpert Object implementation was designed for flexibility and for handling of a large number of rules. The Pascal program was a solution to get some preliminary results and indications of the potentials of those spatial parameters that seemed most useful, within a reasonable amount of time. The main problem with the Nexpert Object implementation was the large number of objects and relations. It was only possible to process very small images. The problem would be reduced if selection criteria could be applied at the retrieval of objects. This would be possible if the objects were stored in a relational database and could be retrieved using SQL statements. It could not be done within this study, as the available version of Nexpert Object (Version 1) could not interface the available relational database management system. Another possible improvement would be to change the processing strategy so that the segments are input and processed one by one instead of all at once. The Nexpert Object based classifier would, however, still be considerably slower than a classifier implemented in a conventional programming language, mainly due to the slow input/output.

Classification of segments based on their spectral and spatial characteristics is *one* method to extract information from an image. However, as all other methods, it solves some problems and introduces others. Land-use classification from satellite data is a highly complex task, that most likely requires that a number of feature extraction algorithms, working on different levels of detail, are combined [e.g., Argialas and Harlow, 1990].

Conclusions

Multispectral classification of high resolution satellite data generally performs poorly for spectrally heterogeneous land-use categories, such as built-up land. This study has shown that the classification accuracy can be improved by additional classification of the segments in the spectral map, based on some of their spatial properties. The classification accuracy for built-up land improved significantly after the application of a few relatively simple rules. However, a more complex set of classification rules is needed to produce a robust and reliable result.

Two factors favor the use of object-oriented and knowledge-based methodologies for segment-based classification. First, these methodologies provide the concepts to represent segments as objects that have certain properties and belong to certain classes, a representation that corresponds well to the way we intuitively think of the segments. Second, when the classification rules grow more complex, the tools for inference become useful as they allow for a flexible and dynamic implementation of the rules.

This study can be regarded as a "proof of concept." It has demonstrated that it is possible to implement a segment-based classification system in an object-oriented expert system shell — Nexpert Object. The main problem was the inability of the system to handle the number of objects involved, even in the classification of a small (512 by 512 pixels) image. This problem can, however, be reduced by allowing filtering through SQL statements when data are input to the system. Thus, the number of segments in the system at any given time can be reduced to what is actually needed for a certain set of rules. Processing time was another major concern. Compared to a simple Pascal implementation, Nexpert Object did not perform very well. However, as a more complex set of rules is used, it is expected that flexibility is more important than processing time. This speaks in favor of the expert system implementation, despite the slow processing.

Acknowledgments

The project has been financially supported by the Swedish National Space Board (Statens Delegation För Rymdverksamhet, DFR), under contracts number 74/89 and 91/90. Data have been provided by the Regional Planning Office in the Stockholm County. The project was carried out within a Ph.D. program at the Department of Photogrammetry, the Royal Institute of Technology (KTH), Stockholm, Sweden. The author wishes to thank Prof. Friedrich Quiel and Docent Hans Hauska for their guidance and support.

References

Argialas, D., and C. Harlow, 1990. Computational Image Interpretation Models: An overview and perspective. *Photogrammetric Engineering & Remote Sensing*, Vol. 56, No. 6. pp 871–886.

Baker, J., S. Briggs, V. Gordon, A. Jones, J. Settle, J. Townsend, and B. Wyatt, 1991. Advances in classification for land cover mapping using SPOT HRV imagery. *International Journal of Remote Sensing*, Vol. 12, No. 5. pp. 1071–1085.

Civco, D., 1989. Knowledge-based land use and land cover mapping. *Proceedings ASPRS/ACSM Annual Convention*, 2-7 April, Baltimore, pp. 276-289.

Genesereth, M., and N. Nilsson, 1988. *Logical Foundations of Artificial Intelligence*. Morgan Kaufmann Publishers Inc., 405 p.

Gonzalez, R., and P. Wintz, 1987. *Digital Image Processing, 2nd Edition*. Addison-Wesley Publishing Company, 503 p.

Goodenough, D., M. Goldberg, G. Plunkett, and J. Zelek, 1987. An Expert System for Remote Sensing. *IEEE Transactions on Geoscience and Remote Sensing*, Vol. GE-25, No. 3, pp 349–359.

Harmon, P., and D. King, 1985. Expert Systems. Artificial Intelligence in Business. Wiley Press, 283 p.

Huang, Z., 1991. *A Data Model for Representation of Raster Data and Vector Data in GIS*. Photogrammetric Reports no. 56 (to be published), Dept. of Photogrammetry, KTH, Stockholm, 9 p.

Johnsson, K., 1990. Generation of Training Statistics Based on Existing Land-Use Information. *Proceedings International Geoscience and Remote Sensing Symposium IGARSS'90*, Washington D.C., 21-24 May, pp. 2383–2386.

———, 1991. *Rule based extraction of spectrally heterogeneous land use categories from SPOT multispectral data*. Licentiate thesis. Department of Photogrammetry, KTH, Stockholm, 77 p.

Johnsson, K., and J. Kanonier, 1991. Knowledge based land-use classification. *Proceedings International Geoscience and Remote Sensing Symposium IGARSS'91*, Espoo, Finland, 3-6 June, pp. 1847–1850.

Lillesand, T., and R. Kiefer, 1987. *Remote Sensing and Image Interpretation, Second Edition*. Johns Wiley & Sons, 721 p.

Mehldau, G., and R. Schowengerdt, 1990. A C-Extension for Rule-Based Image Classification Systems. *Photogrammetric Engineering & Remote Sensing*, Vol. 56, No. 6, pp. 887–892.

Nagao, M., and T. Matsuyama, 1980. *A Structural Analysis of Complex Aerial Photographs*. Plenum Press, New York. 199 p.

Neuron Data, 1990. *Nexpert Object - Users Manual 1/2*. Neuron Data Inc., Palo Alto, California.

Patterson, D., 1990. Introduction to Artificial Intelligence and Expert Systems. Prentice Hall International, 448 p.

Sadler, G., 1991. Information extraction from remotely sensed images for urban land use analysis. *Proceedings of the European Conference on Geographical Information Systems (EGIS-91)*, Brussels, Belgium, 2-5 April, pp. 955–964.

Westervelt, J. (editor), 1988. *Users and Programmers Manual for the Geographical Resources Analysis Support System Version 3.0*, U.S. Army Construction Engineering Research Laboratory, Champaign, Illinois.

Winston, P., 1984. *Artificial Intelligence, Second Edition*. Addison-Wesley, 524 p.

(Received 6 July 1992; revised and accepted 1 December 1992; revised 19 January 1993)