Remote Sensing in Transition, Goossens (ed.)

Extraction of land use / land cover - related information from very high resolution data in urban and suburban areas

T. Van de Voorde, W. De Genst & F. Canters
Centre for Cartography and GIS, Vrije Universiteit Brussel (VUB), Brussels, Belgium

N. Stephenne & E. Wolff
IGEAT, Université Libre de Bruxelles (ULB), Brussels, Belgium

M. Binard
SURFACES, Université de Liège (ULG), Liège, Belgium

Keywords: VHR satellite images, urban remote sensing, land-cover classification

ABSTRACT: Very High Resolution (VHR) satellite images offer a great potential for the extraction of land-use and land-cover related information for urban areas. The available techniques are diverse and need to be further examined before operational use is possible. In this paper we applied two pixel-by-pixel classification techniques and the object-oriented image analysis approach (eCognition) for a land-cover classification of a Quickbird image of a study area in the northern part of the city of Ghent (Belgium). Only small differences in overall Kappa were noted between the best results of the pixel-based approach (neural network classification with Haralick texture measures) and the object-oriented classification (eCognition). A rule-based procedure using ancillary information on elevation derived from a digital surface model was applied on the pixel-based land-cover classification in order to obtain information on the spatial distribution of buildings and artificial surfaces.

1 INTRODUCTION

Urban decision makers are confronted with a complex and dynamic environment. To be able to conduct a policy aimed towards sustainable regional development, they require up-to-date information, supplied by efficient data-extraction systems that support their decision making processes. Satellite images offer this potential, but until recently their spatial resolution was too coarse to allow them to be operationally used for applications where greater spatial detail is needed, such as urban areas. Very High Resolution (VHR) satellite images implicitly contain a rich source of useful information for urban managers and planners. With two operational VHR satellites in orbit (Ikonos and Quickbird) we have possibilities like never before to obtain information at the urban level for extended areas. The possible applications are wide-ranging: inventorying built-up parcels at a regional scale, impervious surface mapping, mapping and assessing urban green areas, identifying urban morphological complexes, determining building density and amount of open spaces in the city for comparative urban research…

The most common procedure to derive thematic maps from remotely sensed imagery is supervised image classification (Foody & McCulloch, 1995). Cities are, however, spectrally and morphologically complex entities. When we look at them from space we observe a huge set of building materials, each having its own characteristics. Allocating a pixel of unknown class membership to a pre-defined (urban) land-cover class based on its spectral properties is not a straightforward task in such environments. As a result of spectral heterogeneity and spatial variance, pixel-by-pixel classifications of urban areas often result in "speckled" classifications, i.e. thematic maps that lack spatial coherence. This is the case for any classifier that operates on individual pixels. Shadows of trees or buildings and sun glint on roofs or windows complicate matters even more.

In order to derive useful thematic maps from VHR satellite images of urban areas, other approaches than the traditional pixel-by-pixel classification are needed. One way to reduce the salt-and-pepper effect in a pixel-based classification is to apply a standard majority filter (Gurney & Townshend, 1983) or a more sophisticated spatial reclassification technique (Barnsley & Barr, 1996; Gong & Howarth, 1992; Wharton, 1982) within a moving window of fixed size. The use of kernel-based approaches, however, has a number of disadvantages, including the difficulty of selecting an optimal kernel size, and the fact that the kernel is an artificial construct that does not refer to the spatial units that occur in the land-cover scene. To avoid the problems related to the use of kernel-based filtering methods, Barr & Barnsley (2000) propose a reflexive mapping proce-

dure that operates on individual regions, i.e. groups of adjacent pixels that are assigned to the same land-cover class by the classifier. The procedure merges regions that fall below an a priori defined, class-specific area threshold with the smallest neighbouring region that exceeds the area threshold. After each re-assignment the region-based topological structure of the image is rebuilt.

An alternative way to look at the problem of structural clutter, instead of allocating individual pixels to a pre-defined land-cover class one-by-one, is to divide the image into regions of similar pixels prior to classification. These so-called segments do not necessarily have any meaning and can be considered as image primitives. Once they are created, they can entirely be assigned to the land-cover classes by any classifier. Many techniques of image segmentation have been developed (cfr. Pal & Pal, 1993). The most common methods to segment a full image are: global thresholding (a survey of these techniques is given by Sahoo *et al.*, 1988), region growing algorithms, watershed segmentation (Wegner *et al.*, 1997) and texture segmentation algorithms (based on spatial frequencies (Hofmann *et al.*, 1998), Markov Random Field models (Panjwani & Healy, 1995), co-occurrence matrices (Haralick & Shapiro, 1985), wavelet coefficients (Salari & Ling, 1995) or fractal indices (Chaudhuri & Sarkar, 1995)). An alternative segmentation method is the so-called probabilistic segmentation (Gorte, 1998). First, an unsupervised per-pixel classifier is applied on multi-spectral image channels. The segments are created after the classification by grouping adjacent pixels that have been allocated to the same class.

While they all address the problem of the salt-and-pepper effect, these segmentation techniques suffer from significant drawbacks, e.g. over and under segmentation or not being useful at all scales, which make them less suitable for urban land-use/land-cover classification. Building on the idea of image primitives as objects, a novel approach to satellite image classification has recently been developed: object-oriented image analysis (Baatz & Schäpe, 1999). Embedded in the commercial software eCognition, the object-oriented processing of image information involves handling image primitives at different scales as objects. These objects are obtained with the multi-resolution segmentation technique (Baatz & Schäpe, 2000) and placed in a hierarchical network. With this approach, the objects not only have spectral properties. Shape, texture, context and information from super or sub-objects are also present. Many successful land-cover or land-use classifications have been obtained with this approach.

Because deriving useful information on urban areas from VHR satellite imagery is clearly a complex matter, five Belgian research laboratories are working together to tackle this problem in the framework of an OSTC funded research project (SPIDER). As part of the project, several strategies to derive land-cover and land-use related information from VHR satellite images are being tested.

In this paper we applied two pixel-by-pixel classification techniques and the object-oriented image analysis approach to map nine land-cover types for a small test zone located north of the city of Ghent. We also applied a post-classification knowledge based strategy, to improve the pixel-by-pixel classification result, using a DSM as an ancillary data source.

2 LEGEND, STUDY AREA, IMAGE DATA AND SAMPLING STRATEGIES

2.1 *Land-use / Land-cover legend*

The land-use and land-cover legend we adopted in this project is a hierarchical scheme with five main urban land-use types (table 1). They are subdivided into 18 level-two land uses, most of which are further subdivided according to intra-class differences in land cover.

Table 1. Land-use / land-cover legend

Land use 1 (LU1)	Land use 2 (LU2)	Land cover (LC)
1. Buildings	1.1 Isolated house	Grey surface
	1.2 Block of houses	Orange/red surface
	1.3 Low building	Green copper
	1.4 High building	Glass or plastic
	1.5 Other	Bare soil
		Water
2. Road and rail network	2.1 Road	Grass
	2.2 Parking	Crops
	2.3 Railway	Shrub and trees
	2.4 Square	Mixed
3. Hydrology	3.1 Water body	
	3.2 Watercourse	
4. Vegetation	4.1 Urban green area	
	4.2 Agriculture	
	4.3 Forest	
5. Miscellaneous	5.1 Sport or recreative area	
	5.2 Graveyard	
	5.3 Construction site	
	5.4 Other	

The definition of the land-cover legend was an iterative process. We started out with a preliminary version that contained more land-cover types than the current one. First classification tests with an Ikonos image of Brussels showed us that we had been too ambitious. A lot of urban surface types with a grey appearance could not be spectrally distinguished. This was for instance the case for concrete, asphalt, cobblestone, slates, etc. All these classes were aggregated into a super-class labeled “grey surfaces”. The “mixed” land-cover class is present for visual interpretation in the context of validation, and was not used for classification purposes.

2.2 *Image data and preprocessing*

The area we selected for this study, measures 643 by 2071 meters, and is situated in the northern part of Ghent (fig. 1). It contains diverse land-cover and land-use types. The industrial buildings next to the canal and the railway infrastructure in the east are directly connected with early twentieth century worker housing near the middle of the area. They contrast with the modern-day residential housing found in the western part.

Figure 1. Quickbird PAN image of the study area © Digital Globe

A fragment of a Quickbird bundled image product (PAN and multi-spectral), taken on August 23rd 2002, and covering the study area was fully 3D ortho-rectified. Nine ground control points were measured using DGPS in real-time mode. The fixed GPS antenna of the Department of Geography of the University of Ghent was used as reference station. A reference digital surface model (DSM) was created using a set of aerial photographs at a 1/12000 scale. The DSM was developed with the photogrammetric software *VirtuoZo* and the actual ortho-rectification was accomplished with IMAGINE OrthoBASE®.

2.3 *Sampling strategies*

To train the classifiers and assess the results of the classifications, training and validation samples of the study area were gathered. This was accomplished by on-screen visual sampling using a PAN-XS fused image and high resolution aerial photographs as a backdrop. Six classes from the land-cover legend were retained: mixed was omitted and bare soil, green copper and crops were not present in the study area. To these classes we added a shadow class to take the omnipresent shadows into account. For each of these seven land-cover types, six pixels were sampled inside 20 polygons, for a total of 120 pixels per class (fig. 2). The sample points were verified on the field where necessary. From the 20 polygons per class, seven polygons were randomly selected as validation data and 13 as training data.

After training and validation data for all classes had been gathered, the "grey surfaces" class was further split into three classes with different tonal levels of grey: light, medium and dark. This was done by applying an ISODATA-cluster analysis on the reflectance values of all pixels within a grey-mask. Three steps were used for defining the mask. First an NDVI image was calculated from the red and near infrared bands. The NDVI image was then thresholded to separate vegetation from other classes. Second, a ratio image was calculated using the red and blue bands. This allowed us to extract all red surface pixels from the non-vegetated pixels identified in the first step. Third, all reflecting surfaces, shadow and water were extracted from the mask resulting from the previous step by excluding all pixels that have consistently high or low values in all spectral bands. Performing an ISODATA-clustering on values within the mask only, ensures that the resulting classes represent different tonal levels of grey and are not influenced by data values that do not belong to grey surface types.

Figure 2. Sampled pixels inside training and validation polygons

We created three different training sets from the polygons selected for training. One contains four pixels, selected within each polygon, that were considered to be spectrally "typical" for the class to which the polygon belongs. Another contains the four typical pixels plus two atypical pixels per polygon. In a third approach, blocks of 3x3 pixels were constructed around each training pixel in order to increase the sample size. Following an inspection of the class signature histograms for each of the three sets, training pixels that were considered to be outliers were removed.

The impact of the type of training set on the performance of the classifiers was tested in the classification stage. The training set that was obtained by constructing the 3x3 blocks around all individually sampled pixels was only used in the per-pixel classification approach. The same validation data were used to compare all strategies.

On a small confidence site within the study area, we performed an exhaustive visual interpretation of land cover (fig. 3) and land use according to the class definitions shown in table 1. This interpretation was also used in the validation phase.

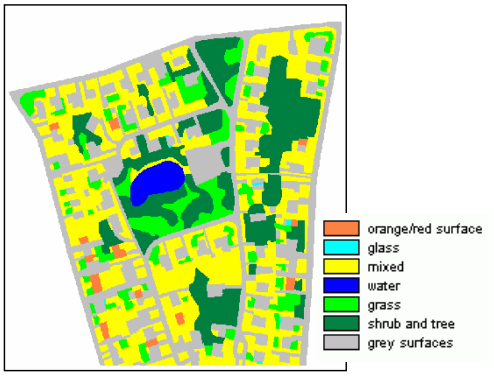

Figure 3. Visual interpretation of land cover for a subset of the study area

3 LAND-COVER CLASSIFICATION

Two major land-cover classification strategies have been applied. The first technique is the conventional pixel-by-pixel classification approach using a supervised classifier. Both maximum likelihood and neural network classifiers have been tested. The other approach is object-oriented image analysis, as implemented in the commercial software *eCognition.*

3.1 *Per-pixel land-cover classification*

The maximum likelihood classifier was applied without a-priori probabilities. The neural networks were built and trained with the Neuralware Predict software. They are characterized by cascade learning (Fahlman & Lebiere, 1990); this means that network architecture is not fixed in advance, but that nodes are iteratively added to the hidden layer during the training process. The effect of adding a node is assessed with an independent set of test cases, randomly selected from the training data. If adding hidden nodes to the network does not improve performance according to the selected error measure (average classification rate in this case), network training is halted.

We used the four multi-spectral channels of the Quickbird image as input for the classifiers. The PAN channel and an NDVI image were also added, in an attempt to improve classification performance. We also trained a neural network with in addition to the spectral channels, texture images based on Haralick co-occurrence matrices (Haralick & Shapiro, 1985), calculated within a moving window on the PAN image. The texture images were based on the angular second moment (ASM), inverse difference matrix (IDM) and entropy (ENT) measures. They were calculated with window sizes five and eleven, for each of the four directions (horizontal, vertical, first and second diagonals). The images calculated for horizontal and vertical directions were also combined to one magnitude image (scalar product of the two directions). The genetic based variable selection of Predict was used to select significant variables for input in the neural network.

Table 2 shows the results of the two per-pixel classifiers, for each training strategy and for the different combinations of input data. The table lists the Kappa indices (Rosenfield & Fitzpatrick-Lins, 1986) calculated on the validation pixels. The "spectral" scenario includes the four multispectral image bands and the derived NDVI band. The "PAN" scenario includes the same bands plus the panchromatic band. The texture scenario includes all bands from "PAN", plus the selection of texture channels chosen by Predict's variable selection. For "atypical" these are ASM in the horizontal direction with window size 11 (ASM-hor-11), ENT-hor-5, IDM-hor-11. For the "typical" training set, only ASM-hor-11 was selected. For "polygon": ASM-1st diagonal-11, ENT-hor-11, IDM-vertical-5, ENT-magnitude-11 and IDM-magnitude-11.

Table 2. Kappa indices for all per-pixel classification scenarios

samples	NN			ML	
	spectral	PAN	texture	spectral	PAN
typical	0.74	0.75	0.74	0.77	0.77
a-typical	0.73	0.79	0.83	0.77	0.79
3x3	0.77	0.81	0.81	0.77	0.78

From the table we see that the strategy that uses a neural network classifier, trained with both typical and a-typical pixels and with texture added to the spectral input channels, provides us with the highest Kappa index (0.83). This classification is shown in figure 4 for a small subset of the study area.

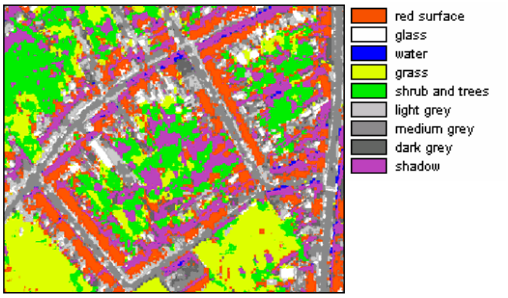

Figure 4. Neural network classification with typical and a-typical training, and with texture added to the spectral input channels

Compared with the ML classifier the NN seems very sensitive to the selected input bands and to the type of training data. Using both typical and atypical training data seems to yield the best results. Adding the PAN image band and texture increases the performance compared with the scenario where only

spectral image information is used. Compared with the "a-typical" training scenario the 3x3 block training does not perform much better. We therefore decided to omit the 3x3 scenario from our future research.

Table 3. Kappas per land-cover class for the best per-pixel classification scenario

Land-cover class	Kappa index
Red surface	0.83
Glass or plastic	0.55
Water	0.85
Grass	0.94
Shrub and trees	0.88
Light grey	0.94
Medium grey	0.81
Dark grey	0.83
Shadow	0.84

The per-class Kappas for the best approach are shown in table 3. Glass appears to cause most problems in the classification. This is probably due to the fact that glass or plastic do not have typical signatures because of the differing incidence angle of the sun. They are often confused with other surface materials .

3.2 *Object-oriented classification*

Multi-resolution image segmentation (Baatz & Schäpe, 2000) is the first step in object-oriented image analysis with the eCognition software. This region growing technique starts from objects with the size of one pixel. Then it subsequently merges adjacent image objects into bigger ones with a procedure that minimizes the weighted heterogeneity criterion of the newly created image objects. The criterion is set by the user before starting the segmentation process. The weight is equal to the object's size. The merging of a region stops when no neighbour can be joined without exceeding the threshold. Because of its higher resolution, in this study the PAN image was chosen to obtain the image primitives. Real image objects have a more identifiable shape and higher spatial contrast in this band.

In a second step, the image primitives that were obtained from the segmentation were classified using the PAN band, the four multi-spectral bands and an NDVI image. We used two training sets: the one with only the typical pixels and the one with both typical and a-typical pixels. The training pixels were overlaid with the segments to obtain training segments for the nearest neighbour classifier embedded in eCognition. In an attempt to improve classification performance we also added extra information in the form of brightness and ratio neo-channels. Brightness of an object is defined as the sum of the mean reflectance values of this object in all bands. The ratio of an object for a certain image band is the object's mean reflectance value divided by the sum of its mean reflectance values in all other bands.

An optimal setting for the heterogeneity criterion (scale factor of 4.7) was chosen by evaluating its impact on classification accuracy. Figure 5 shows the effect of the scale factor on classification performance. Other criterion settings (colour, shape, smoothness, compactness) were left at their default levels.

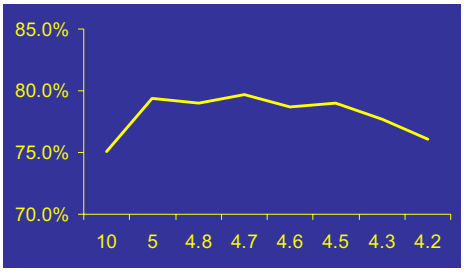

Figure 5. Effect of scale parameter size (x-axis) on classification performance (Kappa on y-axis).

The results of the classifications with this optimal segmentation level are listed in table 4.

Table 4. Kappas for eCognition scenarios

samples	eCognition		
	spectral	PAN	Brightness + ratio
typical	0.71	0.79	0.79
a-typical	0.71	0.80	0.82

Just as for the per-pixel approach, using only spectral data in the classification yields the worst results. Adding the PAN band causes a steep increase of Kappa. Adding additional information like brightness and ratio calculations results in only slightly better classifications. Figure 6 shows the best result of the eCognition approach. From table 5 we learn that the glass or plastic class causes the same problems as for the per-pixel based approach.

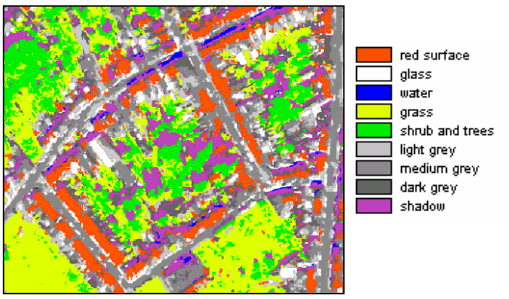

Figure 6. eCognition classification with typical and a-typical training, and with brightness and ratio calculations added to the spectral input channels

Table 5. Kappas per land-cover class for the best eCognition classification scenario

Land-cover class	Kappa index
Red surface	0.88
Glass or plastic	0.58
Water	0.88
Grass	0.85
Shrub and trees	0.81
Light grey	0.97
Medium grey	0.81
Dark grey	0.84
Shadow	0.78

4 FROM LAND COVER TO LAND USE

Thematic maps obtained by a classification of urban surface types can in most cases not be directly used for urban and regional planning or management. Planners or decision makers typically need maps that show urban objects (e.g. houses). They are usually less interested in the constituent elements of these objects (e.g. red surfaces, grey surfaces). Therefore, we need techniques to derive object-level information from the land-cover classifications. We also need to deal with the inherent problems of a land-cover classification, i.e. shadows, and especially for pixel-by-pixel classifications, with the "salt-and-pepper effect".

As a first test to derive meaningful information from land-cover classification results, we applied a rule based procedure on the best neural network land-cover classification of the study area. In this test we mainly focused on the identification of buildings and artificial surfaces.

First, we added information on elevation, derived from the digital surface model (DSM) that was created for ortho-rectification. With an empirically determined threshold of six meters we split each of the nine land-cover classes in two altitude levels: ground level (< 6m) and building level (>6m). The threshold was kept uniform for the entire area because of the sparse relief. After the split we also aggregated the three grey-level classes into one grey class. This theoretically results in a total of 14 land-cover classes.

After this preparatory work, we applied 9 spatial knowledge-based rules in sequence. The rules work on pixels that define a region, i.e. adjacent groups of pixels that belong to the same class.

Let $\mathbf{C} = \{c_1, c_2, c_3,\ldots c_l\}$ be the set of all possible land-cover types, L_1 and L_2 two mapping functions, defined by the neural network classifier, that assign each pixel to the class that corresponds to the highest activated node and to the second highest activated node respectively, and P a post-classification mapping function that re-defines a pixel's class label based on knowledge-based rules.

Moreover, let $\mathbf{A} = \{a_1, a_2,\ldots a_n\} | \forall a \in \mathbf{A}$: $L_1(a) = c_k$ and $\mathbf{B} = \{b_1, b_2,\ldots b_m\} | \forall b \in \mathbf{B}$: $L_1(b) = c_l$ be two adjacent regions with all n (m) pixels allocated to classes c_k and c_l respectively, which we can shortly write as follows: $L_1(\mathbf{A}) = c_k$ and $L_1(\mathbf{B}) = c_l$.

Using these definitions, the rules that have been applied are of two forms:

IF $L_1(\mathbf{A}) = c_k$ AND $L_1(\mathbf{B}) = c_l$ THEN $\forall a \in \mathbf{A}$: $P(a) = c_l$

IF $L_1(\mathbf{A}) = c_k$ AND $L_1(\mathbf{B}) = c_l$ THEN $\forall a \in \mathbf{A}$: IF $L_2(a) = c_l$ THEN $P(a) = c_l$ ELSE $P(a) = c_k$

Figure 4 shows an example of each type of rule. The first rule in this example is used to merge glass and reflecting surfaces above 6 meters with adjacent regions of the grey roof type. The second rule reduces the amount of shadow on regions that actually are shrub. A similar rule is applied to remove shadows from grey surfaces at ground level.

We also applied a structural filter to clean up the map at intermediate steps between the sequential executions of the rules. This filter uses an area threshold of 16 pixels to assign all smaller regions to the largest neighbouring region.

(1) IF region = "glass at building level" AND neighbouring region = "grey class at building level" THEN add all pixels in region to "grey class at building level"

(2) IF region = "shadow at ground level" AND neighbouring region = "shrub at ground level" THEN add pixel in region to "shrub at ground level" IF second most prominent class for the pixel is shrub

Figure 7. Examples of knowledge-based rules

Applying the rules and the filter significantly improves the structure of the classified image and allows us to extract buildings as well as artificial surfaces at ground level. We were also able to remove a large part of the shadows, as can be seen from a comparison between fig. 8 and fig. 9.

By overlaying the identified roofs with the building outlines from the visual interpretation (fig. 10), we can visually assess the quality of the result. Most buildings are well identified, although some small ones are not. Also, the shape and extent of the buildings is not always correct. This is mostly due to errors in the DSM near the borders of the buildings. The buildings that are identified in the eastern part of fig. 10 are located outside the visually interpreted area.

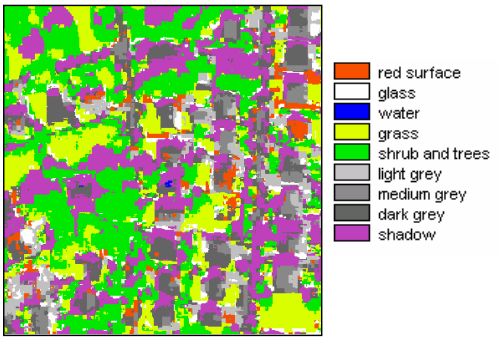

Figure 8. Original neural network classification

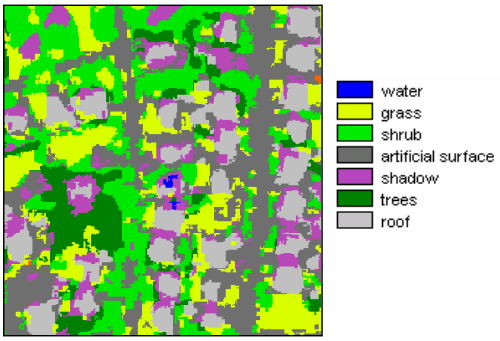

Figure 9. Classification result after applying knowledge-based rules and structural filtering

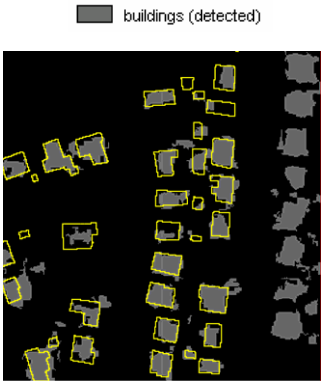

Figure 10. Overlay of building edges from the visual interpretation with roofs detected with the rule-based procedure.

In the future we will refine the technique that is proposed in order to enable and improve the extraction of different types of urban objects.

5 CONCLUSIONS AND FUTURE WORK

From the tests we did we cannot conclude unambiguously that one type of classifier is preferable over the other. The Kappa indices of the best per-pixel scenario and the best object-oriented scenario are not significantly different. Also visually, we cannot unequivocally state which classifier performs best. Some classes appear more structured using eCognition, others do not.

In the near future we will refine the procedures we used here by applying them on four other test zones in Ghent and on study areas in the cities of Brussels and Liège. For these two cities other methods will need to be developed to insert information on the height of objects into the rule-based procedure because of the more pronounced relief there. Depending on the outcome of an in-depth user survey we also plan to develop EO-based applications aimed at fulfilling specific user needs.

ACKNOWLEDGEMENTS

The Belgian Federal Office for Scientific, Technical and Cultural Affairs is gratefully acknowledged for funding the work that is presented in this paper (project SR/00/02). The authors also wish to thank Dennis Devriendt from the Geography Department of the University of Ghent for his work on DSM building.

REFERENCES

Barnsley, M.J. &. Barr, S.L. 1996. Inferring urban land use from satellite sensor images using kernel-based spatial reclassification. *Photogrammetric Engineering and Remote Sensing*, 62(8), pp. 949-958.

Baatz, M. & Schäpe, A. 1999. Object-oriented and multi-scale image analysis in semantic networks. *Proc. of the 2nd International Symposium on Operationalization of Remote Sensing,* August 16th-20th 1999. Enschede. ITC.

Baatz, M. & Schäpe, A. 2000. Multiresolution segmentation – an optimization approach for high quality multi-scale image segmentation. In Strobl, Blaschke & Greisebener (eds), *Angewandte Geographische Informationsverarbeitung XI.* Beiträge zum AGIT-Symposium, Salzburg 1999.

Barr, S.L. & Barnsley, M.J. 2000. Reducing structural clutter in land cover classifications of high spatial resolution remotely-sensed images for urban land use mapping. *Computers and Geosciences*, 26, pp. 433-449.

Chaudhuri, B. & Sarkar, N. 1995. Texture segmentation using fractal dimension. *IEEE Transactions on Pattern Analysis and Machine Intelligence,* 17(1), pp. 72-77.

Fahlman, S.E. & Lebiere, C. 1990. The cascade-correlation learning algorithm. *Tech. Rep. CMU-CS-90-100*, Carnegie Mellon Univ., Pittsburgh, PA.

Foody, G.M. & McCulloch, M.B. 1995. The effect of training set size and composition on artificial neural network classification. *International Journal of Remote Sensing*, 16, pp. 1707-1723.

Gong, P. & Howarth, P.J., 1992. Land-use classification of SPOT HRV data using a cover-frequency method. *International Journal of Remote Sensing*, 13(8), pp. 1459-1471.

Gorte, B. 1998. Probabilistic Segmentation of Remotely Sensed Images. *ITC Publication Series,* 63.

Gurney, C.M. & Townshend, J.R.G. 1983. The use of contextual information in the classification of remotely sensed data. *Photogrammetric Engineering and Remote Sensing*, 49(1), pp. 55-64.

Haralick, R.M. & Shapiro, L.G. 1985. Image segmentation techniques. *Computer Vision, Graphics, and Image Processing*, 29(1), pp. 100-132.

Hofmann, T., Puzicha, J. & Buhmann, J. 1998. Unsupervised texture segmentation in a deterministic annealing framework. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(8), pp. 803-818.

Pal, N.R. & Pal, S.K. 1993. A review on image segmentation techniques. *Pattern Recognition*, 26(9), pp. 1277 – 1294.

Panjwani, D. & Healey, G. 1995. Markov random field models for unsupervised segmentation of textured color images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(10), pp. 939-954.

Rosenfield, G.H. & Fitzpatrick-Lins, K. 1986. A coefficient of agreement as a measure of the thematic classification accuracy. *Photogrammetric Engineering and Remote Sensing*, 52(2), pp. 223-227.

Sahoo, P.K., Soltani, S. & Wang, A.K.C. 1988. A survey of tresholding techniques. *Computer Vision, Graphics, and Image Processing*, 41(2), pp. 233-260.

Salari, E. & Ling, Z. 1995. Texture segmentation using hierarchical wavelet decomposition. *Pattern Recognition*, 28 (12), pp. 1819-1824.

Wegner, S., Oswald, H. & Fleck, E. 1997. Segmentierung mit der Wasserscheidentransformation. *Spektrum der Wissenschaft* 6, pp. 113-115.

Wharton, S.W. 1982. A contextual classification method for recognizing land-use patterns in high-resolution remotely sensed data. *Pattern Recognition*, 15(4), pp. 317-324.