AUTOMATED LAND-USE MAP PRODUCTION FROM AERIAL PHOTOGRAPHS OF PUERTO RICO USING A CONTEXTUAL RECLASSIFICATION ALGORITHM

Johannes van der Kwast, Josefien Delrue, Luc Bertels and Guy Engelen

Flemish Institute for Technological Research (VITO)
Boeretang 200, BE-2400, Mol, Belgium
hans.vanderkwast@vito.be, josefien.delrue@vito.be, luc.bertels@vito.be, guy.engelen@vito.be
http://www.vito.be

Commission WG IV/4

KEY WORDS: land-use change modelling, land-use mapping, contextual reclassification, aerial photography

ABSTRACT:

The Xplorah Planning Support Systems has been under development since 2000 to provide the Puerto Rico Planning Board with an advanced capacity to forecast land-use changes as the result of various scenarios and to assess alternative planning and policy options in their fully integrated, dynamic and spatial context. To fully deploy the analytic capabilities of Xplorah, high quality geographically referenced data are a necessity. These are typically available from Geographical Information Systems at the Planning Board, or other government agencies. The most crucial of all GIS data layers are high quality time series of land-use maps. In order to facilitate regular updating of the land-use map featuring all 21 land-use classes of Xplorah, while assuring a high accuracy, an automatic image processing procedure is being developed. In this procedure the contextual reclassification algorithm OSPARK (Optimised Spatial Reclassification Kernel) is applied to level-sliced, segmented pre-classification of digital aerial photographs with blue, green, red and near-infrared bands. Like SPARK, it is a kernel-based reclassification algorithm, however it automatically adapts the kernel size as a function of spatial variation in the neighbourhood of each pixel to be classified. The algorithm uses the configuration of objects within a kernel in order to effectively discriminate functional land-use classes and land cover. This paper presents the automatic classification procedure and evaluates the results using independent land-use data derived from visual interpretation. Suggestions to improve the procedure are given.

1 INTRODUCTION

With a population just below 4 million living on an area less than 10,000 km^2, the Commonwealth of Puerto Rico is a densely populated island state. The island is unique thanks to its vibrant cultural life, the architecture of its towns, the remnants of its past, and its natural resources: coral reefs, beaches, caves, and tropical forests. It experiences many of the problems typifying small island nations: an open and relatively small economy, a strong dependency on externally provided fossil energy, an overall rugged landscape with a concentration of the population and activities in the coastal zone, fragmentation and loss of the limited high quality agricultural land, high pressure on the coastal wetlands, deforestation, flooding, pollution of surface waters, scarcity of drinking water, solid waste, etc. Moreover, Puerto Rico is located on the path of many a tropical storm entering the Caribbean seas. Some of these gain hurricane force and cause tremendous damage.

Perceptive of the policy and planning problems thus posed, the Graduate School for Planning (GSP) of the University of Puerto Rico (UPR) initiated in 1999 the development of Xplorah, an analytical instrument supporting integrated spatial planning. Since then it has become the planning tool of the Puerto Rico Planning Board to design planning measures and assess their effectiveness. Xplorah is a constrained cellular automata based land-use model of the kind developed by White, Engelen and Uljee (White et al., 1997). It consists essentially of coupled models representing processes operating at three hierarchically embedded geographical levels: the island, the 78 municipalities and the cellular level consisting of cells covering 5.76 ha. At each level, a model representation is adapted to the precise needs of the problems studied, the degree of sophistication required and spatial resolution applied. The models simulate changes in all variables on a yearly basis from the present until a future state, typically 2030. At the cellular level, the land-use functions modelled are: natural vegetation, rangeland, forest, agriculture, construction, mining, industry, trade and services (high and low density) forest reserves, mangroves and swamps. In addition the model comprises land uses that are treated as static features in the model: sea, beach, fresh water, public and recreational areas, utilities, infrastructure, rocky cliffs and shelves.

To fully deploy the analytic capabilities of Xplorah, high quality geographically referenced data are a necessity. The most crucial of all GIS data layers are time series of land-use maps. Xplorah is currently upgraded from its 240 m to a 60 m resolution, enabling effective support for the development of the Puerto Rico Land Use Plan and for addressing practical planning and zoning problems at the level of the island (Isla) and the municipalities (Municipio).

In order to facilitate regular updating of the land-use map featuring all 21 land-use classes of Xplorah at a resolution of 60 m, while assuring a high accuracy, an automatic image processing procedure is being developed. The procedure proposed in this study uses a contextual reclassification algorithm applied to a preliminary classification of digital aerial photographs with blue, green, red and near-infrared bands. The proposed method is tested for a part of the municipality of Mayagüez, located at the west coast of Puerto Rico.

2 METHODS

2.1 The Optimised SPARK algorithm

The concept of contextual reclassifiers is based on the idea that information captured in neighbouring cells or information about

patterns surrounding the pixel of interest may provide useful supplementary information in the classification process. A strong relationship between the spatial structure of urban areas and its functional characteristics have been demonstrated in previous studies (**?**). The Optimised SPARK (OSPARK) algorithm is based on the Spatial Reclassification Kernel (SPARK (Barnsley and Barr, 1996)), which was extended to automatically adapt the kernel size to the spatial variation detected around the pixel to be classified. SPARK is a contextual reclassification method that examines the local, spatial patterns of land cover in a square kernel or moving window. The centre pixel of the kernel is classified according to the arrangement of adjacent pixels. The SPARK method consists of three phases (Barnsley and Barr, 1996):

1. Producing a land-cover map using any type of pixel-based spectral classifier from a remotely sensed image, further referred to as the 'initial land-cover map';
2. Defining SPARK decision rules based on local, spatial patterns of land cover in typical land-use types;
3. Reclassifying the initial land-cover map into land-use types, based on the decision rules of phase 2.

For the production of the initial land-cover map in the first phase, several methods can be used, such as unsupervised clustering techniques, supervised classification techniques, or sub-pixel classification methods (Van der Kwast et al., in review).

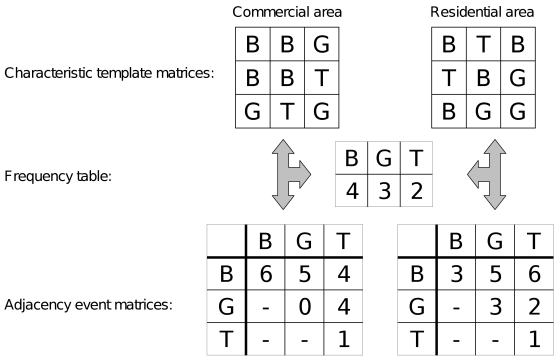

Figure 1: Simulated 3 by 3 kernels of different urban land-use types with corresponding frequency table and adjacency event matrix. Only the upper triangle is considered because $\mathbf{M}_{ij} = \mathbf{M}_{ji}$. B = building, G = grass and T = trees

The decision rules of the second phase are inferred from template kernels that are representative for the land-use classes to be derived. Template kernels are quantified by template ($\mathbf{T}_k$) matrices. One or more $\mathbf{T}_k$-matrices can be used to characterise each land-use type allowing subtle differences in spatial arrangement of pixels in the initial land-cover map to be accounted for (Barnsley and Barr, 1996). Figure 1 shows examples of template matrices. Template matrices are so-called adjacency event matrices that are calculated by counting the frequency of the pixel-based classes positioned next to each other as well as diagonally within each template kernel. Each pair of neighbouring pixels is called an adjacency event. It should be noted that the two template kernels in Figure 1 cannot be discriminated using only class frequency and ignoring the spatial configuration.

In the third phase of the SPARK algorithm, for each pixel in the initial land-cover map an adjacency event matrix is produced by means of a moving window. These adjacency event matrices are called $\mathbf{M}$-matrices.

A similarity index, Δ_k, is calculated to compare the $\mathbf{M}$-matrices derived from the initial land-cover map with all $\mathbf{T}_k$ matrices. Δ_k is calculated using the following equation (Barnsley and Barr, 1996):

$$\Delta_k = 1 - \sqrt{0.5 \cdot N^{-2} \cdot \sum_{i=1}^{c} \sum_{j=1}^{c} \left(\mathbf{M}_{ij} - \mathbf{T}_{k_{ij}}\right)^2} \quad (1)$$

where $\mathbf{M}_{ij}$ is the adjacency event in a c by c matrix, $\mathbf{T}_{k_{ij}}$ is a template matrix for land-use class k, N is the total number of adjacency events in the kernel and c is the number of classes in the per-pixel classified input map. Δ_k can range from 0 to 1. If Δ_k equals 0, $\mathbf{M}_{ij}$ is completely different from $\mathbf{T}_k$, while a value of 1 means that they are identical.

In the final stage, each pixel in the input land-cover map is assigned to the land-use class of the $\mathbf{T}_k$-matrix for which it has the highest Δ_k-value, resulting in a land-use map and a similarity map. The similarity map shows for each pixel the maximum value of the similarity index Δ_k. The procedure is summarized in the shaded part of the flowchart in Figure 2.

The most important factors that determine the quality of the SPARK classification are the representativeness of the $\mathbf{T}_k$-matrices and the choice of an optimal kernel size. The optimal kernel size is a function of the resolution of the images and the scale of spatial variation in the initial land-cover map (Barnsley and Barr, 1996). Kernels that are too large in respect of the land-use objects will increase the effect of edges, while too small kernels possibly do not include all spatial variation. Because of these considerations, some extensions to the original SPARK algorithm have been developed resulting in the OSPARK algorithm (Van der Kwast et al., in review).

Figure 2 shows the flowchart of the OSPARK algorithm. OSPARK iteratively applies the SPARK algorithm for kernel sizes with an apothem from 1 to W pixels. The term apothem is defined as the distance from the centre pixel to a side of a square kernel. After W iterations an integration operator analyses the resulting stack of W similarity maps and assigns the class that corresponds with the optimal Δ_k-value for each pixel.

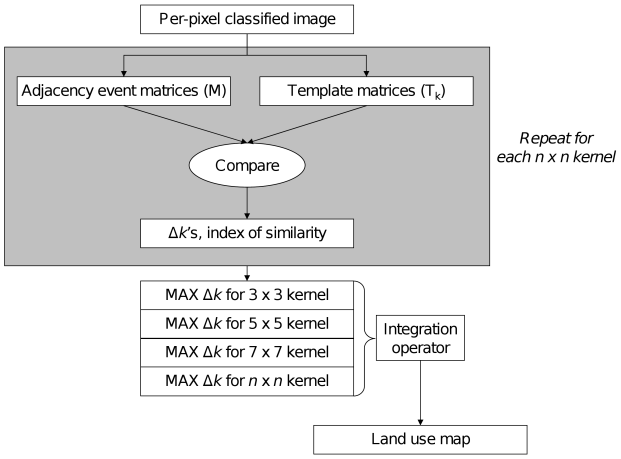

Figure 2: Flowchart of the OSPARK algorithm. The shaded part shows the original SPARK algorithm that is iterated for a range of kernel sizes in the OSPARK algorithm

Since larger template matrices tend to be less unique for a particular land-use type (Sluiter et al., 2004), the Δ_k-value eventually converges to 1 with increasing kernel size. For this reason, the

maximum Δ_k-value for a pixel is not the optimal value. The integration operator of the OSPARK algorithm determines the optimal Δ_k-value for each pixel depending on two possible cases for the evolution of Δ_k with increasing kernel size. The integration operator first checks if local maxima exist. In the case local maxima are present, the first local maximum above a user-defined minimum Δ_k-threshold value is determined and the corresponding land-use class is assigned. In the case that local maxima are absent, the curve converges to $\Delta_k \approx 1$ and the integration operator assigns the class to the pixel when the Δ_k-value changes less than 0.05 and is higher than the threshold value. The user-defined minimum Δ_k-threshold value is used to prevent classification with a low Δ_k-value. If the Δ_k-value of a pixel remains below the minimum Δ_k-threshold value through all iterations, it will be assigned a missing value. The outputs of the algorithm consist of the derived land-use map, a map containing the similarity value corresponding to the optimal kernel size for each pixel.

2.2 Remote sensing data

During the months of October, November and December 2009, thousands of multispectral images were acquired over Puerto Rico, using the ADS40 SH52 digital image sensor of Fugro Earthdata, Inc. Each frame covers 10K by 10K pixels in four spectral bands (Red, Green, Blue and NIR). Flying at an altitude of 2900 m, a ground resolution of 0.3 m was obtained. The onboard positioning system consisted of an Inertial Navigation System (INS) combined with a Differential Geographic Positioning System (DGPS). The INS provides accurate attitude parameters (roll, pitch, yaw), whereas the DGPS records the aircraft's altitude and position. Both are needed for accurate geometric correction of the images. Histogram matching was applied during image pre-processing in order to ensure that all images have a comparable reflectance.

2.3 Decision tree classification

The initial land-cover map needed as input for reclassification with the OSPARK algorithm was derived from the remote sensing images by applying a Decision Tree classification to the images. The Decision Tree classifier is an ENVI® integrated unsupervised classification method. It performs multi-stage classifications by using a series of binary decisions to place pixels into classes. Each decision divides the pixels in an image into two classes based on an expression. Subsequently, each new class can be subdivided into two more classes based on succeeding expressions. This can be repeated as many times as needed. The expressions are operations performed on the blue (B), green (G), red (R) and Near Infra-Red (NIR) bands of the digicam images. These operations were empirically defined resulting in the discrimination of 25 classes for the Puerto Rico images following the procedure described below. In order to investigate the optimal resolution for the OSPARK classification, the initial land-cover map was generated at multiple resolutions between 0.3 m and 15.0 m. An initial land-cover map with a resolution much smaller than the average size of objects will contain too much noise for the OSPARK classification, while a resolution larger than the average size of objects will result in unclear spatial patterns within the OSPARK moving window.

2.3.1 Detection of vegetation Vegetation is detected by using the high reflectance value in the NIR-band which is typical to vegetation. If the pixel value in the NIR-band is greater than in the other three bands, vegetation is assumed (Node 1 in Table 1 and Figure 3). Subsequently, different shades of vegetation are distinguished by using a threshold on the pixel value in the green band (Nodes 2-2, 3-3, 4-5). Table 1 and Figure 3 show the decision rules for vegetation classification.

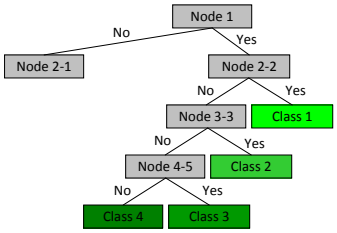

Figure 3: Decision tree for detection of vegetation

Nr.	Class	Node	Rule
-	Vegetation	1	(NIR > B) & (NIR > G) & (NIR > R)
1	Bright vegetation	2-2	G > 125
2	Medium vegetation	3-3	G > 85
3	Dark vegetation	4-5	G > 50
4	Shadowed vegetation	4-5	G ≤ 50

Table 1: Decision rules for classification of vegetation

2.3.2 Detection of shadow and water Low reflectance values in the R, G and B bands are typical for shadow. A threshold is used to separate the shadow pixels (Node 2-1 in Table 2 and Figure 4). The low reflectance values in the NIR band above water bodies is used for the detection of water. However, some dark substrates on land might have low reflectance values in the NIR band as well (Node 3-1). To distinguish those from water, additional constraints were necessary. A threshold on the normalized difference between the NIR and green reflectance value was used to separate water from dark substrates (Node 4-2).

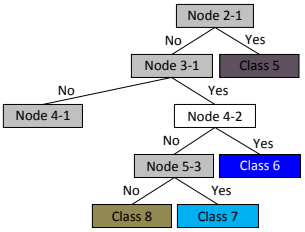

Figure 4: Decision tree for detection of shadow and water

Nr.	Class	Node	Rule
5	Shadow	2-1	(B < 60) & (G < 60) & (R < 60)
-	Water and dark substrates	3-1	NIR < 35
6	Water	4-2	(NIR - G)/ (NIR + G)
7	Class 7	5-3	(G > 60) & (B < G)
8	Class 8	5-3	Not ((G > 60) & (B < G))

Table 2: Decision rules for classification of shadow and water

Unfortunately, some water and dark substrate pixels could hardly be separated. Most of the pixels in Class 7 are water pixels but due to spectral similarity, dark substrate pixels might be present as well (Node 5-3). Most of the remaining pixels, present in Class 8, are dark substrate (Node 5-3). Table 2 and Figure 4 show the decision rules for the classification of shadow and water into four classes.

2.3.3 Detection of bright surfaces and red and orange shades

Bright colours are easily detected by checking the values in the R, G and B bands (Node 4-1). A distinction is made between very bright and extreme bright surfaces (Node 5-2). Table 3 and Figure 5 show the decision rules for the classification of bright surfaces into two classes (Class 9 and 10).

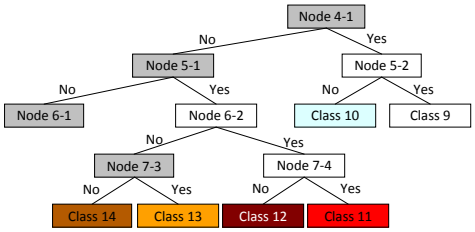

Figure 5: Decision tree for detection of bright surfaces and red and orange shades

Nr.	Class	Node	Rule
-	Bright and extreme bright surfaces	4-1	(R > 200) & (G > 200) & (B < 200)
9	Extreme bright	5-2	B > 230
10	Bright	5-2	B ≤ 230

Table 3: Decision rules for classification of bright surfaces

The red and orange shades in the image are detected by the higher reflectance value in the red band, taking into account a certain threshold value in order to distinguish them from shades of grey (Node 4-1). Next the red and orange shades are separated (Node 5-1). Subsequently, for both distinction is made between the bright and dark shades (Node 7.3 and 7.4). Table 4 and Figure 5 show the decision rules for the classification of red and orange shades into four classes.

Nr.	Class	Node	Rule
-	Red and orange shades	5-1	(R > G + 20) & (R > B + 20)
-	Red	6-2	R > (G + 50)
-	Orange	6-2	R ≤ (G + 50)
11	Bright red	7-4	R > 120
12	Dark red	7-4	R ≤ 120
13	Bright orange	7-3	R > 120
14	Dark orange	7-3	R ≤ 120

Table 4: Decision rules for classification of red and orange shades

2.3.4 Detection of objects with shades of grey

The detection of objects with shades of grey is done by a threshold on the R, G and B bands. In case of shades of grey, all three bands have similar reflectance values (Node 6-1). Different shades are distinguished by a threshold value (Nodes 7-2, 8-3 and 9-5). Figure 6 and Table 5 show the decision rules for the classification of shades of grey into four classes.

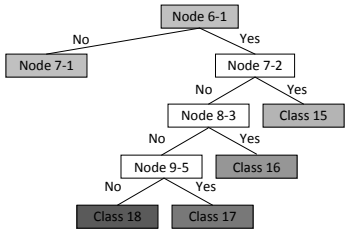

Figure 6: Decision tree for detection of shades of grey

Nr.	Class	Node	Rule
-	Shades of grey	6-1	(B > G - 15) & (B < G + 15) & (G > R - 15) & (G < R + 15)
15	Bright grey	7-2	G > 180
16	Medium grey	8-3	G > 130
17	Dark grey	9-5	G > 90
18	Extreme dark grey	9-5	G ≤ 90

Table 5: Decision rules for classification of shades of grey

2.3.5 Detection of yellow and brown objects

Yellow colours are detected when the red and green band have similar reflectance values and higher than the reflectance of the blue band (Node 7.1). Different shades are distinguished by a threshold value (Node 8.2 and 9.3). Figure 7 and Table 6 show the decision rules for the classification of yellow and brown objects.

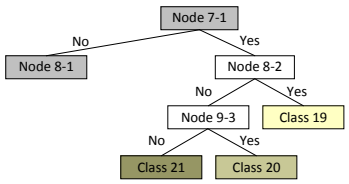

Figure 7: Decision tree for detection of yellow and brown objects

Nr.	Class	Node	Rule
-	Yellow and brown	7-1	(R < G + 50) & (G > B + 10)
19	Bright yellow	8-2	G > 170
20	Medium yellow	9-3	G > 110
21	Brown	9-3	G ≤ 110

Table 6: Decision rules for classification of yellow and brown objects

2.3.6 Detection of objects with remaining colours

The remaining blue colours are detected when the reflectance value in the blue band is higher than those for green and red. The same accounts for the remaining green colours, but here the reflectance in the green band are higher than those in the blue and red band. The remaining purples are separated in bright and dark by a threshold value on the blue band. Figure 8 and Table 7 show the decision rules for the classification of the remaining objects.

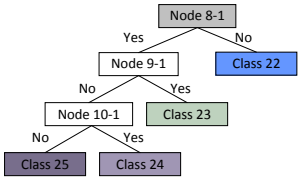

Figure 8: Decision tree for detection of the remaining colours

Nr.	Class	Node	Rule
22	Blue	8-1	(B > G) & (B > R)
23	Green	9-1	(G > B) & (G > R)
24	Bright purple	10-1	B > 110
25	Dark purple	10-1	B ≤ 110

Table 7: Decision rules for classification of remaining colours

2.4 Contextual classification

First, reference data for training and evaluation of the OSPARK algorithm was derived from the Xplorah 2003 land-use map (Román and Castro, 2008), which has a resolution of 15 m and a reported overall accuracy of 97.47%. A subset of 81 ha covering part of the municipality of Mayagüez, was used for testing the methodology. This subset of the land-use map was carefully selected in order to include the most important land-use classes involved in urban dynamics, while limiting its extent for convenient calculation times of the OSPARK algorithm. The centre coordinates of the template kernels were derived by stratified random sampling of 50 points within each class of the land-use map. Van der Kwast et al. (in review) assume that this sampling method results in representative template kernels describing the spatial variability of the different classes. The same procedure was followed to derive an independent set of pixels for evaluation of the contextual classification.

Next, for each resolution of the initial land-cover map the OSPARK algorithm was applied with kernel apothems varying from 1 to 30 pixels. A Δ_k-threshold value of 0.7 was used. The resulting land-use maps were evaluated at their original resolution, i.e. the resolution of the initial land-cover map, and at 60 m resolution by deriving statistics from error matrices.

3 RESULTS AND DISCUSSION

Figure 9 shows the cumulative distribution of cluster sizes in the initial land-cover map. The chart is calculated from the initial land-cover map of 0.9 m resolution and shows an inflection point at 6.0 m. The figure shows that few objects have a size less than 6 m. Since the OSPARK classification infers land-use classes based on the configuration of objects within a kernel, a resolution of the initial land-cover map greater than or equal to 6.0 m seems appropriate in order to reduce the effect of noise at higher resolutions and to speed-up the algorithm.

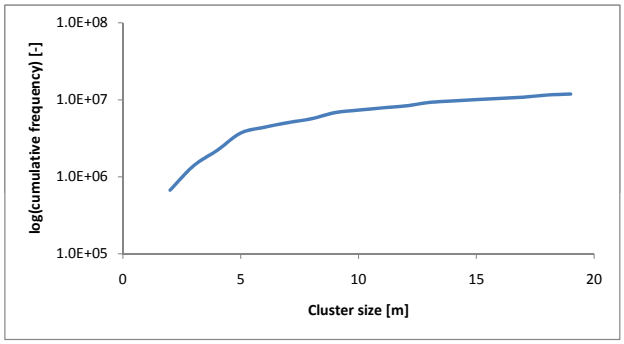

Figure 9: Cumulative distribution of cluster sizes in the initial land-cover map at a resolution of 0.90 m

Figure 9 does not allow the determination of the maximum resolution of the initial land-cover map. In order to investigate this and to confirm the minimum resolution, the OSPARK algorithm has been applied to different resolutions of the initial land-cover map. Figure 3 shows four scenarios of the evolution of the overall accuracy and Kappa when the resolution of the initial land-cover map is decreased. Figure 10(a) shows the evolution of the accuracy when all classes are used and evaluated at the resolution of the initial land-cover map. Figure 10(c) shows the evolution of the accuracy when the classification result is resampled to 60 m and evaluated with the reference map at 60 m. The objective of this study was to develop a methodology for the automatic mapping of land-use at a resolution of 60 m with a high accuracy, containing all classes of the Xplorah land-use change model. Figure 10(c) shows that with the current procedure OSPARK does not produce land-use maps with an overall accuracy higher than 50%. The highest accuracy was obtained with an initial land-cover map with a resolution of 9.0 m. Table 8 shows the user's and producer's accuracy for the individual classes. The classes industry, low-density trade and services, water resources and utilities could be discriminated, while other classes were confused. Merging some of the confused classes resulted in an overall accuracy of 69% and higher class accuracies (Table 9), although residential areas and infrastructure are still hard to classify.

Land use	User	Producer
Forest	0.40	0.45
Agriculture	0.41	0.57
Industry	0.79	1.00
High-density trade and services	0.50	0.25
Low-density trade and services	0.70	0.76
High-density residential	0.35	0.36
Low-density residential	0.55	0.47
Mangroves and swamps	0.37	0.78
Water resources	0.65	0.88
Public and recreation	0.53	0.57
Utilities	0.90	0.84
Infrastructure	0.35	0.39
Rangelands	0.31	0.10

Table 8: User's and producer's accuracy of the OSPARK classification, using all classes and a resolution of 60 m. Classes that do not occur in the subset have been omitted. Overall accuracy: 0.52, Kappa: 0.48

Land use	User	Producer
Non-urban	0.77	1.00
Industrial	0.79	1.00
Trade and services	0.76	0.54
Residential	0.45	0.28
Water resources	0.68	0.88
Public and recreation	0.59	0.57
Utilities	0.90	0.84
Infrastructure	0.35	0.39

Table 9: User's and producer's accuracy of the OSPARK classification, using merged classes and a resolution of 60 m. Overall accuracy: 0.69, Kappa: 0.62

An important factor for the low accuracies for the OSPARK classification using all land-use classes probably lies in the fact that a reference land-use map of 2003 is used, while the imagery was acquired in 2010. The stratified random sample of reference pixels might not be representative for the current land use. Since a more recent land-use map is not available, manual selection of representative reference areas might improve the OSPARK classification.

SPARK and OSPARK have been mostly applied to medium res-

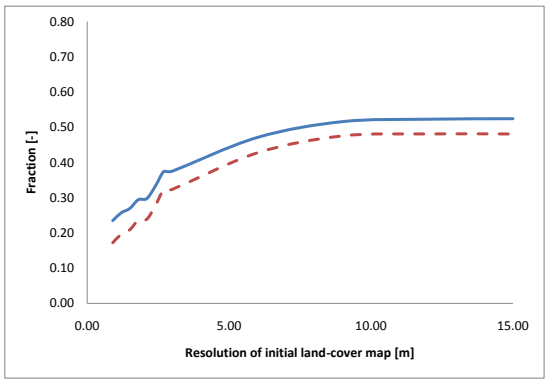

(a) All classes – Original resolution

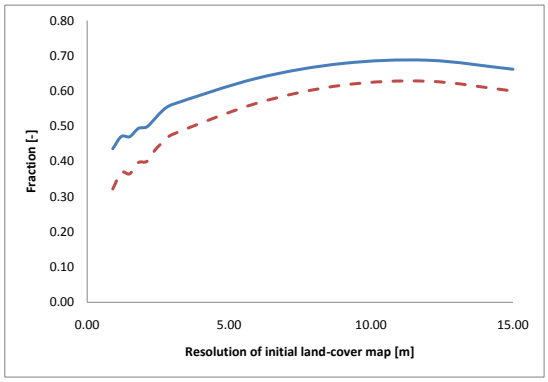

(b) Merged classes – Original resolution

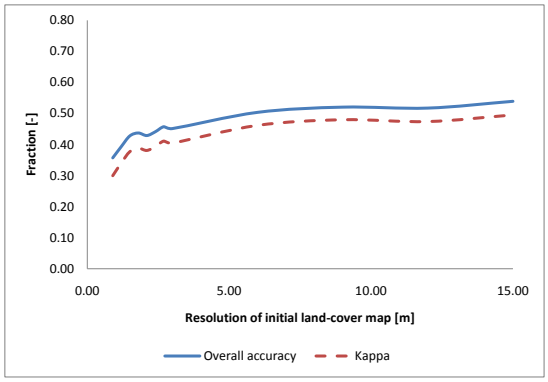

(c) All classes – 60 m

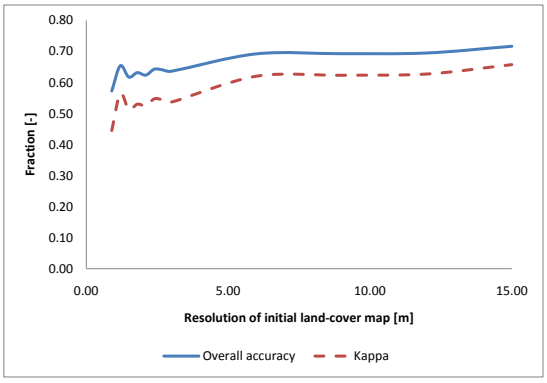

(d) Merged classes – 60 m

Figure 10: Overall accuracy and Kappa for four scenarios

olution (MR) remote sensing images. The advantage of applying OSPARK to high resolution (HR) imagery, such as aerial photographs, IKONOS and QuickBird, is that elementary objects, such as houses, can be identified in the initial land-cover map at a lower level of spatial abstraction. Objects derived from MR images are more generalized and abstract, thus harder to be used by the contextual algorithm A major drawback of increasing the resolution of the initial land-cover maps is the increase in calculation time. The structure of the OSPARK algorithm, however, easily allows reduction of computing times by distributing the iteration process over multiple cores, processors or computers and to centrally collect the results of the iterations for processing with the integration operator (Van der Kwast et al., in review).

4 CONCLUSIONS AND FUTURE WORK

This study proposes an automatic land-use classification procedure based on the contextual classification of high resolution multispectral digital airborne images of the island of Puerto Rico. Results show that an optimal resolution for the initial land-cover map, used as an input for the contextual classification, is approximately 9.0 meters, which corresponds well with the cluster size distribution found in the study area. The use of this optimal resolution prevents large computing times with redundant data at higher resolutions, while preserving the spatial organisation of objects needed in the OSPARK classification. Results of the OSPARK classification show that at currently the derivation of 21 land-use classes is not feasible with an overall accuracy higher than 0.52. Merging classes results in nine important categories that can be discriminated with an overall accuracy of 0.69. The use of more representative reference locations that do not change in time might improve the classification results. Further research should confirm this. Future work should also investigate the feasibility of using the OSPARK algorithm in parallel computing, which is necessary to reduce calculation times if the proposed procedure is applied to the entire island of Puerto Rico.

ACKNOWLEDGEMENTS

The research presented in this paper is funded by the Graduate School of Planning / University of Puerto Rico in the frame of the Xplorah project. The land-use data of the municipality of Mayagüez were made available by GMT Corp.

REFERENCES

Barnsley, M. J. and Barr, S. L., 1996. Inferring urban land use from satellite sensor images using kernel-based analysis and classification. *Photogrammetric Engineering and Remote Sensing* 62, pp. 949–958.

Román, G. and Castro, A., 2008. *Generation of land use maps required for the implementation phase of a spatial decision support system for Puerto Rico: Xplorah 2003 Land Use Map-Technical Documentation*. Technical report, Geographic Mapping Technologies Corporation, San Juan, Puerto Rico.

Sluiter, R., de Jong, S. M., van der Kwast, H. and Walstra, J., 2004. A contextual approach to classify Mediterranean heterogeneous vegetation using the spatial reclassification kernel (SPARK) and DAIS7915 imagery. In: *Remote sensing image analysis: including the spatial domain*. Kluwer Academic Publishers, pp. 291–310.

Van der Kwast, J., Van de Voorde, T., Canters, F., Uljee, I. and Engelen, G., in review. Land-use classification using the optimised spatial reclassification kernel (OSPARK). *International Journal of Applied Earth Observation and Geoinformation*.

White, R., Engelen, G. and Uljee, I., 1997. The use of constrained cellular automata for high-resolution modelling of urban land use dynamics. *Environment and Planning B* 24, pp. 323–343.