METHODOLOGY FOR A
CANADIAN AGRICULTURAL LAND COVER CLASSIFICATION

T. Fisette*, R. Chenier, M. Maloley, PY. Gasser, T. Huffman, L. White, R. Ogston, A. Elgarawany

Agriculture and Agri-Food Canada, 960 Carling Avenue, Ottawa, Ontario, K1A 0C6, Canada -
(fisettet, chenierr, maloleyr, gasserpy, huffmant, whitelo, ogstonr, elgarawanya)@agr.gc.ca

ABSTRACT:

The National Land and Water Information Service (NLWIS), an Agriculture and Agri-Food Canada (AAFC) initiative, is tasked with mapping 'circa 2000' land cover for agricultural areas of Canada with an overall accuracy of greater than 85%. The method for mapping land cover is being developed using a Decision Tree (DT) classifier (see5 software) with specific focus on the identification of 3 agricultural classes: grassland, annual and perennial crops. Class separation is based on rule sets automatically derived from ground information. Using an eCognition segmentation on the Landsat data, a per-object majority filter is performed on the output classification. A unique mosaic procedure has been developed using segmented objects and accuracy results from See5. Preliminary results show the AAFC methodology outperforms both the standard Maximum Likelihood Classifier (MLC) and the object based eCognition method. Implementing multi-date imagery to separate cropland from hay and pasture improves the accuracy of separating other classes, therefore leading to an overall improvement of the classification accuracy. The effect of ancillary data on the classification is dependent upon landscape characteristics. The DT approach, when combined with image segmentation, has proven to meet or overcome our accuracy requirements with minimum analyst intervention.

RÉSUMÉ:

Par l'entremise du Service National d'Information sur les Terres et les Eaux (SNITE), Agriculture et Agroalimentaire Canada (AAC) entreprend la cartographie de la couverture du sol pour le secteur agricole canadien. Celle-ci sera réalisée pour l'année 2000 avec une précision anticipée supérieure à 85%. Une méthode par arbre de décision (logiciel see5) est utilisée en mettant l'emphase sur l'identification de 3 classes agricoles: prairies, cultures annuelles et cultures pérennes. Un ensemble de règles dérivé automatiquement de sites d'entraînement servira à différencier les différentes classes de couverture du sol. Un filtre de majorité est appliqué sur la classification résultante en utilisant les objets issus d'une segmentation des données Landsat (logiciel eCognition). Les résultats de la segmentation ainsi que la précision mesurée par le logiciel see5 permettent de gérer efficacement le processus de mosaïque. Les résultats préliminaires de classification démontrent que la méthode par arbre de décision surpasse la méthode par maximum de vraisemblance et l'approche orienté objets (eCognition). L'utilisation de données multi dates améliore la séparation entre les cultures annuelles et pérennes. Du même coup, la précision globale de la classification s'en trouve nettement améliorée. L'effet d'introduire des données auxiliaires dans le processus de classification est très variable et dépend des caractéristiques du paysage. L'approche par arbre de décision suivie d'une segmentation d'image nous permet d'atteindre, voire même dépasser nos objectifs de précision tout en minimisant l'intervention de l'analyste.

1. INRODUCTION

Increasing public pressure for environmental preservation, continued emphasis on high outputs in agricultural production and competition from other land uses are resulting in higher demands for careful management of soil, water and air resources. Scientific support for land use planning must start with reliable information relating to current resource conditions. One of Canada's most glaring gaps in that aspect is the lack of current, accurate, complete and accessible digital land cover information. Current efforts to assess and improve production sustainability rely on a collection of agricultural land cover products that are out of date and of varying scales, classification schemes and accuracy levels.

Preparation of regional land management tools such as nutrient management plans, environmental farm plans, best management practices, risk management plans and emergency measures plans require reliable data on the location of production activities with respect to waterways, topography, soil capability and cultural features such as roads and settlements. For purposes such as these, land use information that is capable of identifying features such as individual fields and farmsteads is required.

Although the priority is for land use (economic) information, land cover information is more easily mapped and can serve as an approximation of land use. Land cover information derived from earth observation data has been identified within the National Land and Water Information Service (NLWIS) as a fundamental dataset to support many NLWIS objectives and related departmental projects, programs and activities.

AAFC is developing a method for medium resolution (30m) land cover mapping of the agricultural area of Canada. The primary goal of the classification is to accurately delineate agricultural lands: this includes annually cultivated (cropland), perennially cultivated (pasture/forages), and grasslands. Other contextual classes (such as developed or forested lands) are also identified, however accuracies may be lower for these classes,

as the classification process has been developed to target the separation of agricultural lands. The classification methodology under development at AAFC integrates a supervised Decision Tree (DT) classifier with an object oriented approach. This paper presents an overview of the methodology and shows initial results from selected test sites.

2. METHODS

2.1 Data Acquisition and Image Pre-Processing

To date, about 80% of the required Landsat imagery has been acquired. Although plans call for the systematic mapping of all agricultural areas, methods are being developed and tested on agro-ecologically representative sites across Canada (Figure 1). These include 3 sites in eastern Canada (Ottawa-Ontario, Victoriaville-Québec and PEI-Prince Edward Island) and 7 sites in western Canada (Minnedosa-Manitoba, Swift Current-Saskatchewan, Meadow Lake-SK, Lethbridge-Alberta, Calgary-AB, Peace River-AB and Vancouver-British Columbia). Each site is comprised of one Landsat satellite scene ~ 175 km x 175 km. The fragmented landscape in eastern Canada, where fields sizes are variable, contrast with the larger fields in the prairies and the mountainous area of the western part of the country.

Figure 1. Test sites location

For each scene, at least 2 'circa 2000' images from early and mid growing season (typically May and July) were acquired. Landsat images from different years are used when two suitable images are not otherwise available in the same growing season. The Landsat data were not atmospherically corrected. Results from other studies (Song *et al.*, 2001; Champagne *et al.*, 2005) demonstrate that an atmospheric correction is unnecessary when classifying multi-temporal data, assuming that the training data comes from the same images as those being classified. All optical images were orthorectified using a 3-D multi-sensor physical model developed at the Canada Centre for Remote Sensing (Toutin, 1995) and implemented in PCI software. Ground control points (GCP) used for the geometric correction came from the CPLIC1 database of Geobase, an initiative overseen by the Canadian Council on Geomatics. A Tasseled Cap (TC) algorithm, in which Bands 1 to 5 and 7 were converted into three channels (brightness, greenness, wetness) was applied to each image pair. The TC transformation reduced

data processing while maintaining accuracy (Kauth and Thomas, 1976; Lillesand and Kieffer, 1994).

Separation of different land cover types can be enhanced using spatial textures (Stuckens *et al.*, 2000). An intra-pixel texture measure was computed from the local variance in an adaptively placed 3 by 3 window of the green band on the mid season image. The location of the adaptively located window used to calculate local variance is the window with the lowest local variance of all windows that include the pixel (Ryherd and Woodcock, 1996). The intent is to calculate texture from the object or region to which a pixel belongs and avoid the edge enhancement associated with the calculation of texture from windows centered on the pixel in question. A semi-automated cloud and shadow masking technique based on a comparison of TC values between clouded and cloud-free images was applied to every image. Also, ancillary data such as DEM, DEM derived products (slope, aspect) and soil information were used as additional input to the classifier.

Training data for the DT classifier was collected from a variety of sources such as crop insurance databases, aerial photography and visual interpretation of Landsat images. A minimum of 500 sample polygons were selected in order to separate the 10 classes outlined in Table 1. The number of sample sites collected per class is proportional to the area covered by that class within a specific scene. Class proportions were estimated from the EarthSat GeoCover LC-1990 product (MDA, 2006) and the 2001 Canadian census of agriculture recompiled at the scene level. Through a systematic random sampling grid, reference data to assess the accuracy of the final land cover product (mosaic) will be collected.

Land Cover	
Unclassified	■
Water	■
Barren	■
Developed	■
Shrubland	■
Wetland	■
Grassland	■
Annual Crops	■
Perennial Crops and Pasture	■
Coniferous	■
Broadleaf	■

Table 1. Land cover classes

2.2 Classification Methodologies

Classification was performed on a scene by scene basis because the dynamic nature of crop rotations, crop growth and harvest patterns create significant reflectance differences between adjacent satellite scenes within the temporal period encompassed by scene availability. All classifications were conducted using digital number (DN) values. Time series images were used for each classification in order to separate annual crops from hay and pasture, as the classes have varying spectral characteristics over the growing season. The DT method, as implemented in see5 software, is a multivariate model based on a set of decision rules defined by combinations of features and a set of linear discriminant functions that are

applied at each test node. Decision boundaries and coefficients for the linear discriminate function are estimated empirically from the training data. The DT method was chosen because of its ability to handle discrete data, processing speed, independence of the distribution of class signatures, interpretable classification rules (Friedl and Brodley, 1997; Pal and Mather, 2003) and cost effectiveness. Advanced options such as pruning and boosting have also been incorporated into the decision tree classification process to improve the accuracy of the algorithm. Figure 2 provides an overview of the full classification process.
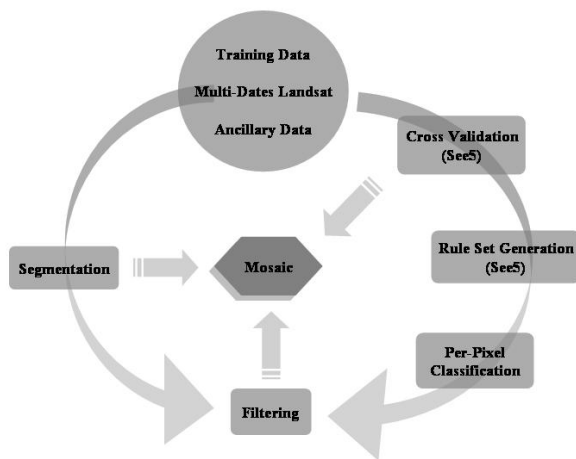

Figure 2. Generalized classification process

A Graphical User Interface (GUI) developed by AAFC and Statistics Canada automatically manages the classification processes through see5, PCI Geomatica and ArcGIS. The training sites, ancillary data and spectral information from the tasseled cap transformation were entered into the See5 software. A cross-validation was performed on the centroids of the training site polygons, providing an estimate of the accuracy: centroids were divided into *n* blocks of roughly the same size and class distribution. For each block in turn, a classifier was constructed from the samples in the remaining blocks and tested on the samples in the hold-out block. Results give the analyst the ability to revisit the training sites or to proceed with the classification. A two-level hierarchical classification scheme (Figure 3) is implemented in the methodology. This hierarchical approach allows the landscape to be segmented into general classes within which more detailed classes can be discriminated.
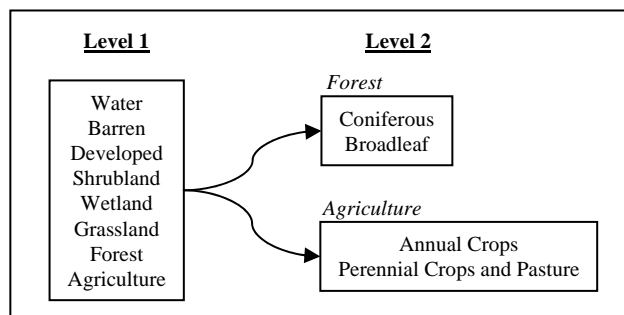

Figure 3. Two-level hierarchical classification scheme

This portion of the methodology has been compared with a conventional Maximum Likelihood Classifier (MLC) on 5 sites. eCognition software classification results were also evaluated at the Lethbridge and Ottawa sites. A 3-level top down approach was built using a combination of Nearest Neighbor and membership functions. Annual and perennial crops were differentiated on the last level at a scale factor of 10 and a spectral weighting of 0.8 and a spatial weighting of 0.2.

2.3 Post-Processing

In order to improve accuracy and aesthetics, an eCognition segmentation and majority filter was applied to the DT classification. However, clouds, overlapping and non-overlapping areas must be processed independently, so a thematic layer partitioning the scene into homogenous zones representing these areas was created. This layer and the multi-temporal Tasseled Cap bands were input to the eCognition segmentation algorithm that derives polygons (object primitives). A spatial scale of 10 was found to be a good compromise between number of objects and individual field representation. The polygons were imported into ESRI, ArcGIS Spatial Analyst and a majority zonal statistic was calculated on the per pixel classification. This assigned the majority class value within each object primitive to the entire polygon.

The filtering impact on the classification was also evaluated using other filters (mode and sieve) on the Ottawa and Lethbridge sites. 3x3 and 7x7 pixel mode filters were used. The sieve filter removes all polygons smaller than a given size and merges them with their largest neighbours. Thresholds of 15 and 30 pixels were tested.

Each zone within the thematic layer and in consequence each individual object, is linked to an accuracy value from the cross-validation step. This accuracy information guides the automated mosaic process. Objects sharing a line segment with an image edge or cloud/shadow edge are eliminated from the scene being merged. Object deletion occurs on the overlapping area only if data exist on the mosaic. A second phase compares cross-validation accuracies of the overlapping area and deletes objects with lower values. The remaining objects are then pasted onto the mosaic product and accuracy values are updated.

3. RESULTS

Classification trials with the DT classifier were performed on ten test sites. However, quantitative accuracy assessment does not support all sites; in some cases, a qualitative evaluation of the classification was appropriate to support the methodology development. Initial results indicate that the DT approach outperformed the MLC method on the 5 sites (Table 2). Overall accuracy differences varied from about 5% to 13%. Large differences could be an indication that training statistics are not suitable as they indicate non-homogeneity within classes and are non-Gaussian. DT methodologies permit the integration of disparate geospatial data and unlike maximum likelihood classifiers, the DT approach does not make any assumptions regarding the statistical distribution of these data (McNairn et al., 2005).

Comparisons between the DT classifier and eCognition (for classification purposes) were also investigated. Overall results for eCognition were almost identical to the MLC method. Preliminary results show eCognition has very good potential in

defining the broad agriculture class (annual and perennial crops) for both the Lethbridge and Ottawa sites. Fields are very well delineated on the per-object classification. However, distinction between annual and perennial crops is less satisfactory, particularly with the Lethbridge scene where fields are a mix of dryland and irrigated land. eCognition requires significant user input and expert knowledge of the area being classified. The potential for higher accuracies exists, but even if attained, the exact process -mix of spectral and spatial settings- cannot be applied to other scenes. In comparison, the DT method is much simpler. Expert knowledge (aside from acquiring training sites) is not needed and the generation of rules is automatic. With DT, classification time is more a function of processor power than human intervention. Due to time (operational) constraints and considering the fact that AAFC needs an approach that is relatively constant across Canada, using eCognition alone was not considered an option.

	DT	MLC
PEI*	86.3%	80.9%
Ottawa	88.3%	83.8%
Meadow Lake*	96.5%	89.5%
Lethbridge	87.7%	74.5%
Vancouver*	95.6%	83.7%

Table 2. DT and MLC overall accuracy assessment comparison (* no differentiation between annual and perennial crops)

Implementing multi-date Landsat imagery to the DT process has lead to an overall improvement in classification accuracy of up to 10 percent. Marked improvements were observed for agriculture classes (up to 20% for the user and producer accuracies). The effect of adding ancillary data to the classification process is dependent upon the landscape (Table 3). In mountainous regions such as Lethbridge and Vancouver, DEM and derived products (slope, aspect) improved overall accuracies by more than 4%. Effects on other sites with flat or low rolling topography were negligible. In all cases, introducing texture slightly improved classification results; Classes with high spectral heterogeneity such as developed areas were better represented. The Canadian Land Inventory (CLI) index, which defines the land capability for agriculture from soils information at the 1:250 000 scale, has proven to be useful for some sites. A slight increase in accuracy was observed for scenes with large portions of land considered as non-suitable for agriculture such as the Canadian shield on the Ottawa scene and the mountainous area around Lethbridge and Vancouver.

	Change in Overall Accuracy		
	Texture	DEM & derived products	CLI
PEI	+0.7%	0.0%	+0.2%
Victoriaville	n/a	-0.7%	-1.8%
Ottawa	+0.2%	+1.0%	+2.1%
Lethbridge	+2.5%	+5.2%	+1.2%
Vancouver	n/a	+4.1%	+0.9%

Table 3. Changes in overall accuracy when adding selected ancillary data to the DT classifier

The object based majority filter yielded superior results to the mode and sieve filters. For these latter two, accuracy increases with the filter size (Table 4). Accuracy results are similar to the object based filter but class edges are rounded and linear features are lost (Figure 4). The object based majority filter added between 2% to 4% to the per-pixel classification overall accuracy. Although very close in accuracy to the mode and sieve filters, it is clearly superior visually. Class edges are enhanced and linear features are better represented. Statistically, an increase in reference data along edges and narrow features such as roads would probably further distinguish the object based filter from the other two.

	Lethbridge	Ottawa
Raw	85.5%	92.5%
Mode 3x3	87.0%	94.0%
Mode 7x7	89.0%	94.5%
Sieve (15 pixels)	88.2%	93.9%
Sieve (30 pixels)	88.7%	94.1%
Object based	89.3%	94.2%

Table 4. Filtering effect on classification accuracy

Merging two contiguous classified scenes without seam lines or rules creates obvious discontinuities (Figure 5). Scene edges and cloud borders - due to variations in input imagery dates - become apparent. Inconsistencies could be easily removed by manually editing seam lines. This would take a considerable amount of time however, and end product accuracy would be compromised.
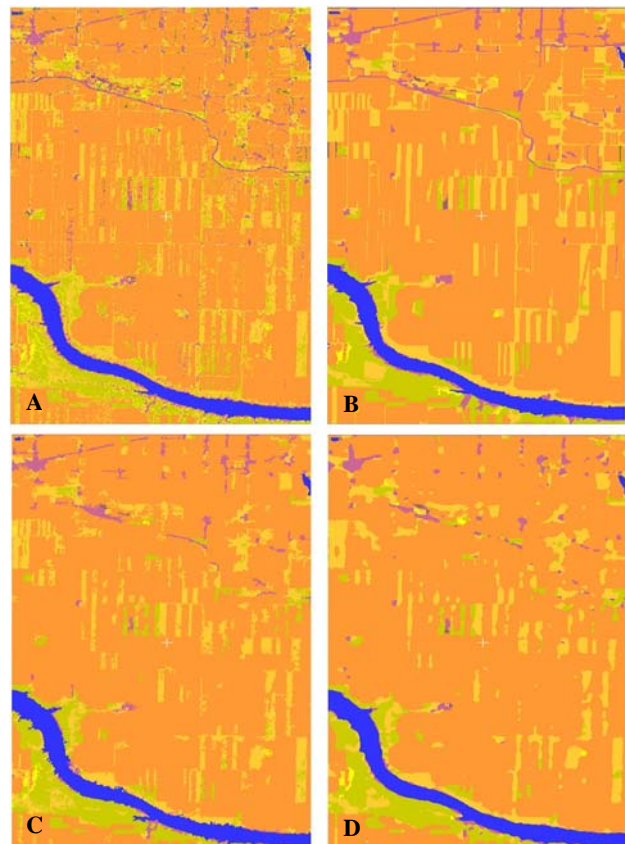

Figure 4. Post-filtering effect: A) Original classification; B) Object based filter; C) 3x3 mode filter; D) Sieve Filter (30 pix.)

The object-based mosaic approach optimizes final product accuracy and minimizes visual inconsistencies by generating cut lines along segmented object edges (Figure 6). However, pasting one object-based classification over another may introduce slivers along the cut lines: objects of the newly classified scene may not match those of the existing mosaic product. Most of these small polygons can be removed during the overlapping process by using a minimum tolerated distance between vertices: vertices that fall within the set cluster tolerance are snapped together. Trials conducted thus far show that remaining clusters have a very small visual impact at the scene scale.
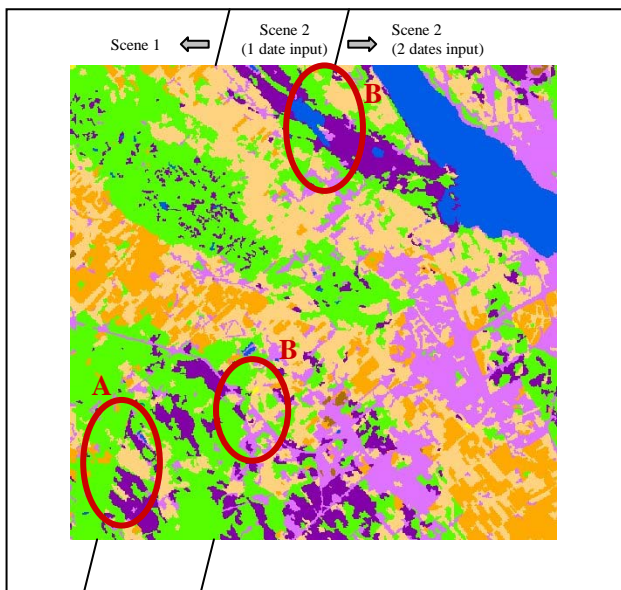

Figure 5. Inconsistencies on the mosaic product: (A) between 2 adjacent scenes and (B) within scene, along multi-date imagery edges (spring and summer images are used to classify portion to the right and spring image is used to classify strip to the left).
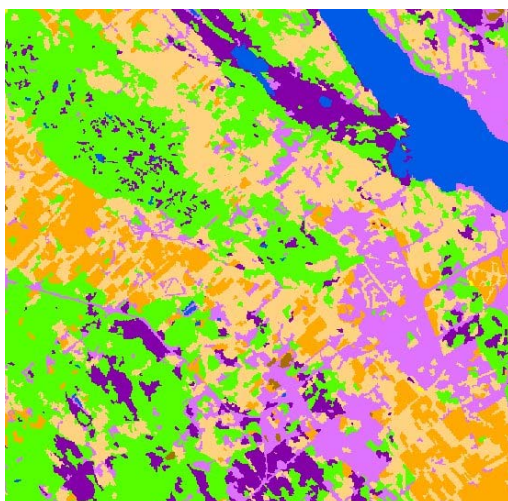

Figure 6. Mosaic output from the object based approach. Inconsistencies are no longer visibles.

4. CONCLUSIONS

A Decision Tree (DT) method has been developed to successfully classify Canadian agricultural lands. Post filtering and the combination of images into mosaics make use of object segmentation. PCI, See5, ESRI, and eCognition software were integrated in to a custom GUI written in VB.Net. The DT classifier was tested on several sites across the country, yielding superior results when compared to a MLC or a "pure" eCognition approach. In addition to being a flexible land cover classification tool, the DT approach permits the integration of disparate data sources. The AAFC methodology takes advantage of the segmentation process of eCognition, which reduces the "speckle" effect found in pixel-based classifiers, relying on a more statistically robust algorithm to automate class separation.

Our 85% target accuracy was reached for every site where a quantitative assessment was done. However, considering Canada's diverse agricultural landscape, testing on more sites is required to establish its robustness. Our intent is not to build a static methodology but rather to constantly improve it as our classification work evolves. Integration of other earth observation imagery and geospatial data within the classifier will continue to be explored.

REFERENCES

Champagne, C., Shang, J., McNairn, H., and Fisette, T., 2005. Exploiting Spectral Variation from Crop Phenology for Agricultural Land-Use Classification. In: Proceedings of SPIE, *Remote Sensing and Modeling of Ecosystems for Sustainability II*, San Diego, California, Vol. 5884, July 31 – August 5, 2005.

Friedl, M.A. and Brodley, C.E., 1997. Decision tree classification of land cover from remotely sensed data. *Remote Sensing of Environment*, 61(3), pp.399-409.

Kauth, R.J. and Thomas, G.S., 1976. The Tasseled Cap - A Graphic Description of the Spectral-Temporal Development of Agricultural Crops as Seen by LANDSAT. In: *Proceedings of the Symposium on Machine Processing of Remotely Sensed Data*, Purdue University, West Lafayette, Indiana, pp.41-57.

Lillesand, T.M. and Kieffer, R.W., 1994. *Remote Sensing and Image Interpretation*. Wiley & Sons, New York, pp. 523-525.

McNairn, H., Shang, J., Champagne, C., Huffman, E., Smith, A. and Fisette, T., 2005. A multi-sensor approach to inventorying agricultural land use. In: *Proceesings of the 31st International Symposium on Remote Sensing of the Environment*, St.Petersburg, Russia. http://www.isprs.org/publications/related/ ISRSE/html/ sections.html (accessed 18 May 2006)

MDA Federal Inc., 2006. EarthSat GeoCover LC Overview. http://www.mdafederal.com/geocover/geocoverlc/gclcoverview (accessed 15 June 2006)

Pal, M. and Mather, P.M., 2003. An assessment of the effectiveness of decision tree methods for land cover classification. *Remote Sensing of Environment*, 86(4), pp.554-565.

Ryherd, S. and Woodcock, C., 1996. Combining spectral and texture data in the segmentation of remotely sensed images.

Photogrammetric Engineering and Remote Sensing,62(2), pp.181-194.

Song, C., Woodcock, C.E., Seto, K.C., Pax Lenney, M. and Macomber, S.A., 2001. Classification and change detection using Landsat TM data: when and how to correct atmospheric effects? *Remote Sensing of Environment*, 75(2), pp.230-244.

Stuckens, J., Coppin, P.R. and Bauer, M.E., 2000. Integrating contextual information with per-pixel classification for improved land-cover classification. *Remote Sensing of Environment*, 71(3), pp.282-296.

Toutin, T., 1995. Multi-source data fusion with an integrated and unified geometric modelling. *EARSeL advances in remote sensing*, 4(2), pp.118-129.