The Effects of Polygon Boundary Pixels on Image Classification Accuracy

Paul Boudewyn*, David Seemann, Mike Wulder and Steen Magnussen
Canadian Forest Service
506 West Burnside Road
Victoria, B.C. V8Z 1M5
Canada

*corresponding author: phone (250) 363-0600 email pboudewyn@pfc.forestry.ca

Abstract

The purpose of this study is to analyze the effects that pixels, located at polygon boundaries, have on classification accuracy. Pixels found along the borders of polygons usually contain mixed spectral information, and can be detrimental to classification accuracy. Discriminant analysis was used to predict land cover classes, found in Canada's National Forest Inventory, from a Landsat TM image. The discriminant criteria were derived on a test area of the image, using buffered and non-buffered polygons as training data, and applied to a validate area of the image. Buffering the polygons had no overall positive or negative effect on classification accuracy. There are non-trivial effects for specific cover types, especially the water categories, but the classification accuracy for most categories changed by less than 10% due to buffering. Overall accuracy is quite low as well, usually less than 50%, which suggests that discriminant analysis may not be suited for predicting National Forest Inventory land cover classes from Landsat TM images.

Introduction

Satellite images are an important information resource for forestry, and have a wide variety of applications. They have many potential uses, especially in forest inventory, because of their large ground coverage, their ability to be analyzed and manipulated with various mathematical and statistical procedures, and their increasing resolution and availability. These unique characteristics can be used to predict forest inventory attributes such as land cover types, tree density classes, and tree mortality. The ongoing challenge for foresters is to use or combine conventional techniques for deriving these attributes in innovative ways, so that they can accurately inventory the forest resource.

The purpose of this study is to analyze the effects that pixels, located at polygon boundaries, have on classification accuracy. To do this, we compared the classification accuracy of land cover classes predicted from classifiers trained on Landsat TM image segments, with and without boundary pixels. Pixels found along the boundaries of polygons usually incorporate spectral or class information from adjacent polygons due to rasterization of vector data (Congalton 1997). As a result, these boundary pixels may be different from the pixels found in the interior of polygons. An analysis by Wulder et al. (1999) tends to support this contention, as they found the spectral contents of boundary pixels to be significantly different from that of interior pixels. Since 52% of the pixels in our study site are boundary pixels (Table 1), they could potentially have a strong effect on classification accuracy.

We used discriminant analysis to predict the land cover classes. This classification technique derives a relationship between a categorical variable and a set of interrelated, quantitative

***Paper presented at: "*Remote Sensing and Spatial Data Integration: Measuring, Monitoring and Modeling*", 22nd Symposium of the Canadian Remote Sensing Society, Victoria, British Columbia, August 20th to 25th, 2000.**

variables. This relationship is then used to predict the category of an entity, based on the associated measurements of the quantitative variables (McLachlan, 1992). Our approach in this study comes from a forest inventory perspective, and uses pixel grey levels of the seven Landsat TM image channels as the quantitative variables, and Canada's National Forest Inventory (NFI) land cover classes as the categorical variable. The relationship between the two sets of variables was derived from a portion of the image, or training dataset, and was then applied to the entire image to classify all the pixels. In order to assess the effects of boundary pixels, the training datasets included both buffered and non-buffered polygons.

Materials and Methods

A 1995 Landsat TM scene (30m by 30m resolution) was obtained for the study site, which includes the southeastern portion of New Brunswick, Canada, along with the forest inventory database for the same location and year (Figure 1).

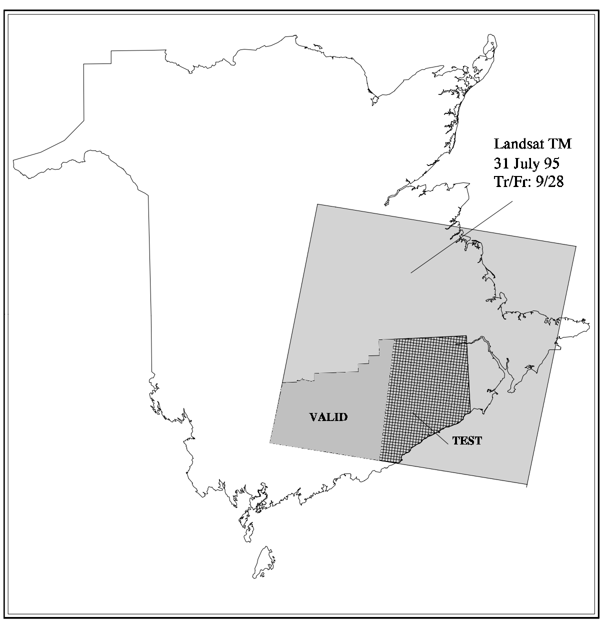

Figure 1. Province of New Brunswick, showing study area, Landsat TM coverage and the test and validate regions.

The image was carefully geo-referenced with the GIS inventory database. A root mean square error of 0.46 pixels was achieved by using 37 ground control points and a first-order polynomial, nearest-neighbour model. This procedure ensures that all the information contained in the forest inventory coverage can be overlaid and matched with the information contained in the Landsat image. We applied a 30 m buffer, or a one-pixel buffer, to all polygons. The two data sources were then merged, and two text files were produced. One for the eastern half of the study area, which we called the test area, and one for the western half of the study area, or the validate area (Figure 1). The following information was retained or added to every pixel: (i) grey level for channels 1 to 7; (ii) National Forest Inventory (NFI) land cover classification for levels 2 and 3; (iii) flag to indicate whether the pixel was found on the boundary (within buffer) or in the interior of a polygon. The NFI land cover classification is a hierarchical system that is used to classify a polygon or area into 5 levels, with level 1 being the broadest class and level 5 being the narrowest (NFI design document, 2000). We classified each pixel into levels 2 and 3 of the NFI scheme, which are as follows: level 2 – land (l), non-treed (n), treed (t) or water (w); level3 – crops/pasture (cp), exposed-land (el), lake (la), ocean (on), river (ri), rock (ro), shrub-low (sl), shrub-tall (st), tree-broadleaved (tb), tree-mixedwood (tm) and tree-coniferous (tc). Level 1 classifies a pixel into vegetated or non-vegetated, while levels 4 and 5 further classify level 1, 2 and 3 into landscape positions and vegetation densities. Levels 2 and 3 of the NFI land cover classes are predicted in this study (Table 1). The discriminant analysis procedure will predict the membership of a pixel into these land cover classes from the 7 Landsat TM channels. We are not going to discuss the statistical theory behind discriminant analysis, but rather, we will outline the steps we used to implement the discriminant procedure using SAS software.

The first step in the analysis was to take the text files created for the test and validate area and normalize each of the 7 TM channels with equation 1.

$$z = \frac{x - \mu}{\sigma} \quad (1)$$

Then, we performed principal component analysis on the normalized values of the TM channels. This helped to reduce the number of variables used in the analysis from 7 channels to 4 principal components, while still explaining almost 98% of the variance. The Princomp Procedure of SAS was used, with the default options. Next, the Discriminant Procedure of SAS was applied to the data from the test area (SAS, 1993). We limited the analysis to polygons that were equal to, or larger than 25 pixels, to produce sufficiently consistent training information. NFI land cover class was set as the class variable and, after much thought, the prior probabilities were set to equal, because we decided that the proportions of land cover found in the test area did not necessarily imply a similar proportion for the validate area. After running the Discriminant Procedure on the test area, it was applied to the validate area. The output, or discriminant functions, derived from the first run on the test area are used to predict a NFI land cover class for each pixel of the validate area. Finally, the entire procedure was repeated several times, once for each NFI land cover class (levels 2 and 3 described previously) and once for each of interior, interior plus boundary, or boundary pixels. Figures 3a and 3b graphically portrays forest polygons and the way they are decomposed for this analysis. The interior (light tone) pixels only are used for the training dataset in one scenario, the interior plus boundary (light plus dark) pixels are used for another, and the boundary (dark) pixels only are used for a third scenario.

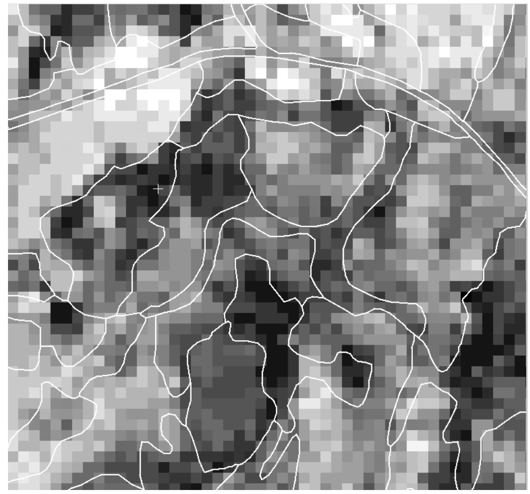

Figure 3a. A magnified portion of the Landsat TM image overlaid with the NFI polygon data.

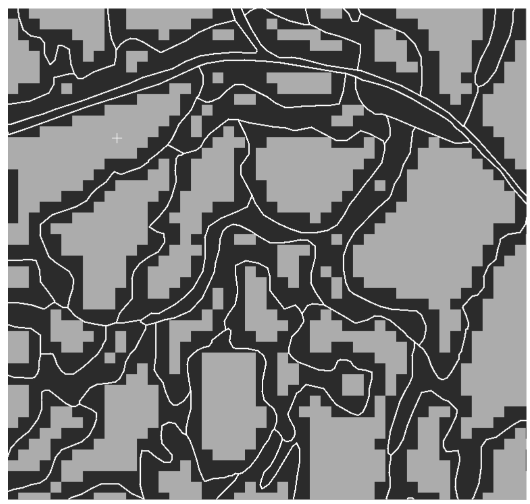

Figure 3b. The same area showing the effects of buffering (bitmap image).

Results and Discussion

Figures 4a, 4b, 5a, and 5b summarize the results in terms of classification accuracy. We define classification accuracy as the proportion of pixels, or polygons, correctly classified into their respective NFI land cover classes. We also use the term buffering interchangeably with removing boundary pixels. Each figure displays the accuracy for the NFI land cover class and for each of the 3 different training datasets used as inputs to the discriminant analysis. The proportions shown in the figures should be used in conjunction with Table 1, which describes the data.

Table 1. NFI land cover classification scheme and descriptive statistics for the test and validate area.

NFI level 2	NFI level 3	Test (training) area					Validate area	
		pixels	polygons	mean polygon size (ha)	% boundary pixels	average posterior probability	pixels	polygons
land – l	exposed-land – el	113330	1213	8.4	77.2	0.246	405912	3928
	rock – ro	3749	14	24.1	30.0	0.546	911	29
nontreed - n	bryoid	0	0	0	0.0	0.0	179	1
	crops/pasture - cp	400372	1571	22.9	32.5	0.396	259414	2557
	shrub-low – sl	142019	1884	6.8	65.0	0.119	349316	5027
	shrub-tall – st	92765	1095	7.6	56.9	0.224	85483	1805
treed – t	tree-broadleaved - tb	990063	9727	9.2	50.9	0.413	826033	11398
	tree-coniferous - tc	1258948	13921	8.1	54.0	0.254	1094526	17605
	tree-mixed – tm	1427575	15561	8.3	54.6	0.206	1670255	22972
water - w	lake – la	8229	76	9.7	42.4	0.732	137810	250
	ocean – on	137	2	6.2	61.3	0.979	702	2
	river – ri	12594	120	9.4	92.3	0.257	174625	203
	Total	4449781	45184	8.9	52.6		5005166	65777

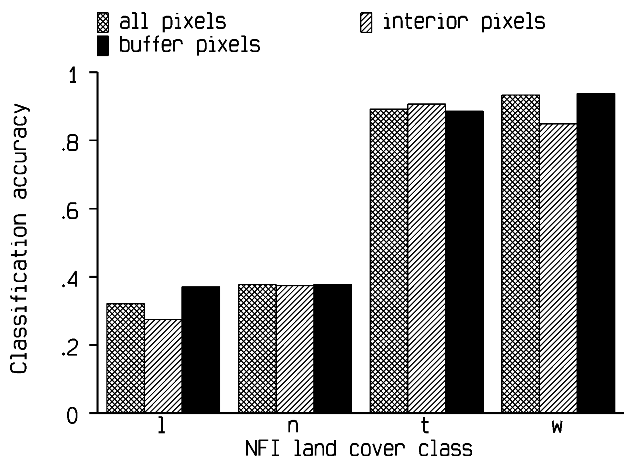

Figure 4a. Classification accuracy of pixels for NFI level 2.

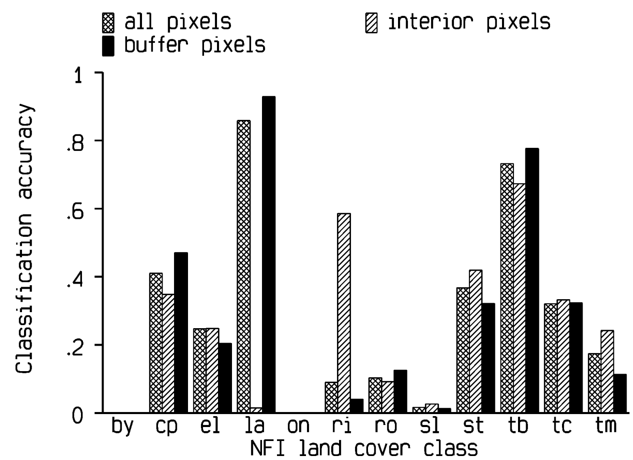

Figure 4b. Classification accuracy of pixels for NFI level 3.

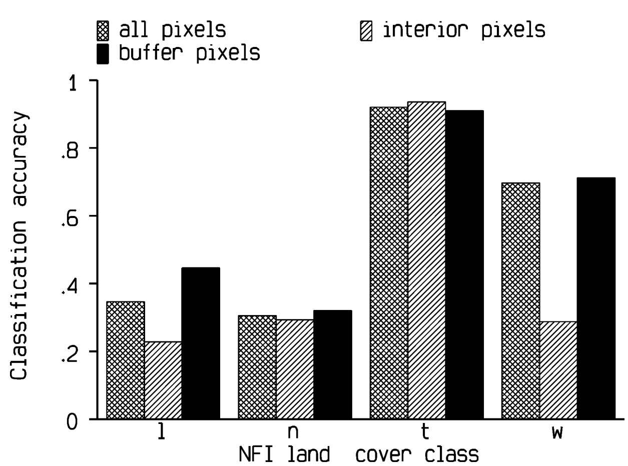

Figure 5a. Classification accuracy of polygons for NFI level 2.

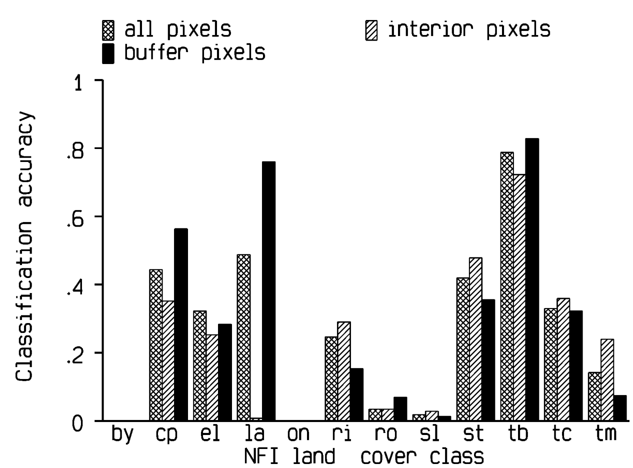

Figure 5b. Classification accuracy of polygons for NFI level 3.

The following general trends are evident from examining these figures, and will be discussed further.

- The classification accuracies follow similar trends for pixels and polygons.
- Overall, there is no consistent positive or negative effect caused by buffering.
- The classification accuracies tend to be quite different among the NFI land cover classes.

The pattern and values of the results shown between Figures 4a and 5a, and between Figures 4b and 5b, are quite similar, which indicates that the results are not affected by the method of summarization. For example, in Figures 4a and 4b, the number of correctly classified pixels is simply totaled for each NFI land cover class over the entire validate area. In contrast, the results in Figures 5a and 5b were calculated by averaging the posterior probability of each pixel by polygon, and labeling the polygon with the NFI land cover class that has the highest average posterior probability. Because the results are similar, it suggests that pixels within a polygon have a relatively low variance in their spectral values, and on average, a clear majority is classified to the same NFI land cover class.

Figures 4a to 5b show that there is no consistent positive or negative effect of using training data with or without boundary pixels. Buffering the polygons does not increase accuracy across all NFI land cover classes, nor does it decrease accuracy. However, there are some interesting trends when we look at the results within specific NFI land cover classes. There is virtually no effect caused by buffering for the non-treed and treed categories in Figures 4a and 5a (<5%), and the same result is evident for the further subdivision of these categories represented in figures 4b and 5b. The accuracy of the crops/pasture, shrub-low, shrub-tall, tree-broadleaved, and tree-coniferous categories all change positively or negatively by less than about 10%. The tree-mixedwood has slightly more of a change at about 15%. These cover classes represent over 85% of the test area, and have average posterior probabilities that are all less than 0.41 (Table1). In other words, there is a large population of pixels, with inherently mixed signatures, which are used for training, and removing boundary pixels obviously has little influence on classification accuracies.

For the land category in Figures 4a and 5a, there is consistent, but small decrease in classification accuracy between 5% and 15% due to buffering. And there is a similar result for the further subdivision of these categories into exposed-land and rock in Figures 5a and 5b. The exposed-land category often occurs as a smaller, more linear shaped polygon, and has over 77% boundary pixels. Losing boundary pixels probably weakens the discriminant functions and reduces the classification accuracy. The strongest effects appear in the water categories. Buffering reduces the classification accuracy by 9% in the water category (Figure 4a) to almost 90% in the lake category (Figure 4b), and increases the accuracy for rivers. The main reasons for these results can be derived from Table 1. There are only 8229 pixels, found in 78 polygons, for lakes in the test area. Buffering reduces this amount to less than 4000, which may not be enough to train on. Even though lake pixels have a relatively high average posterior probability (0.73), which indicates a unique signature, there are just too few pixels remaining to derive good discriminant criteria. Interestingly, most of the remaining lake pixels were predicted as rivers or treed-coniferous, which suggests that they resemble rivers or the dark bands of coniferous trees. It is uncertain why the accuracy for rivers increased from buffering, as they tend to be found in long, narrow polygons. Perhaps they have a unique reflectance that is captured in the analysis, or contain boundary pixels that are very different from the interior pixels.

The overall classification accuracies vary considerably among the NFI land cover classes. A majority of the classes have accuracies less than 50 %, which we consider quite low. This may suggest that the 30 m resolution and spectral characteristics of Landsat TM images are not conducive for delineating the NFI land cover classes. Wynne et al. (2000) report similar findings in their study. They state that is difficult to exceed 85 % accuracy on a per-pixel basis for just two classes, forest and nonforest, and that accuracy is significantly reduced for more specific forest types such as deciduous and coniferous. Also, discriminant analysis may not be the most suitable method for classifying the information content of satellite data. Magnussen et al. (2000) found that other techniques, including fuzzy classifiers, neural nets, and Bayesian analysis, can predict certain land cover classes more accurately.

We have inferred several times that polygon size may influence the accuracy of our predictions. Figures 6a, 6b, 6c, and 6d shows the effects of polygon size on classification accuracy for the crops/pasture, lake, tree-broadleaved, and treed-coniferous NFI land cover classes. Again, the results are somewhat mixed and specific to the cover type, but generally, larger polygons are predicted with as good or slightly better accuracy, than small polygons.

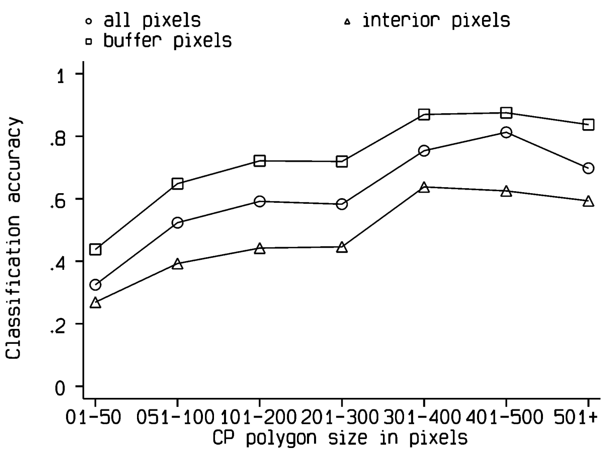

Figure 6a. Classification accuracy vs. polygon size for crops/pasture.

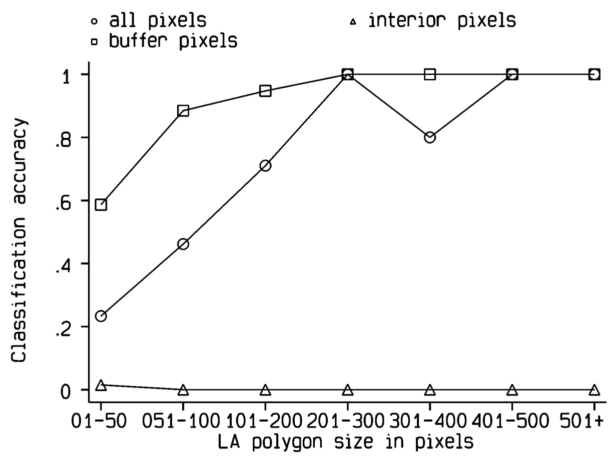

Figure 6b. Classification accuracy vs. polygon size for lakes.

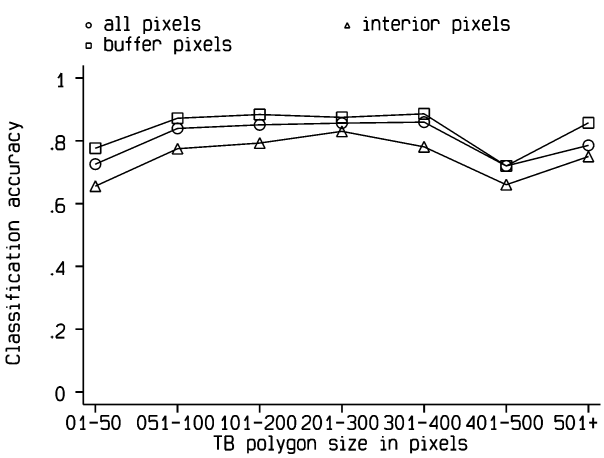

Figure 6c. Classification accuracy vs. polygon size for tree-broadleaved.

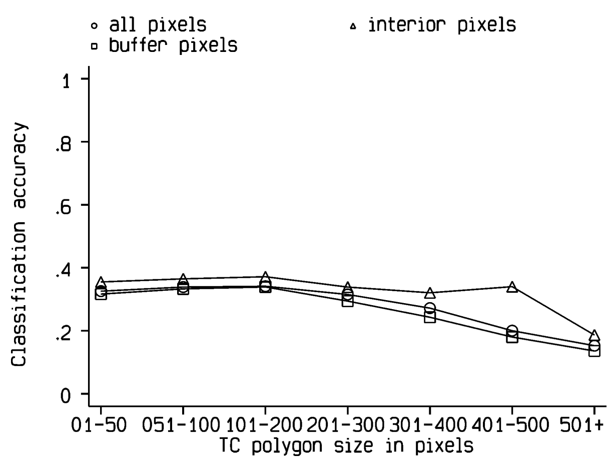

Figure 6d. Classification accuracy vs. polygon size for tree-coniferous.

These figures also show how polygon size interacts with buffering. For the tree-broadleaved and tree-coniferous categories, the effects of buffering are very similar across all polygon sizes. For the crops/pasture category, it appears that the smaller polygons are less affected by buffering than the larger polygons. It is difficult to state any trends for the lake category.

Conclusions

Buffering the polygons intuitively seems like a reasonable method to increase classification accuracy. However, as this study shows, removing the boundary pixels has really no overall positive or negative effect. Also, the classification accuracy is less than 50% for many of the categories, especially for NFI level 3, which we consider quite low. Discriminant analysis does not seem to be an appropriate classification method for predicting NFI land cover classes from Landsat TM images, and our attempt to improve the results with buffering was unsuccessful. There are other factors that have invariably influenced our results, including accuracy of georeferencing and co-registering Landsat TM images with other data sources, the lack of consistent spectral information within polygons of the same label, and time differences between acquisition of inventory and image data. Minimizing these factors is crucial, but we believe that significant improvements in the results may be achieved by using finer resolution imagery and classification techniques more suited to the information content of satellite images.

Acknowledgements

We would like to thank the New Brunswick Department of Natural Resources and Energy for supplying their forest inventory database for this study, and for providing valuable assistance during the database creation and interpretation phase of our work.

References

Congalton, R., 1997. Exploring and Evaluating the Consequences of Vector-to-Raster and Raster-to-Vector Conversion. *Photogrammetric Engineering and Remote Sensing*, 63(4):425-434.

Magnussen, S., P. Boudewyn, M. Wulder, and D. Seemann, 1999. Predictions of forest inventory cover-types using Landsat TM. Submitted to *Silva Fennica, 1999*.

McLachlan, G. J., 1992. Discriminant Analysis and Statistical Pattern Recognition. John Wiley and Sons, Inc., New York.

Natural Resources Canada, 2000. A Plot-based National Forest Inventory Design for Canada. Canadian Forest Service, Pacific Forestry Centre, Victoria, B.C., 140 p. (file report).

SAS Institute Inc., 1993. SAS/STAT User's Guide, Version 6, Fourth Edition, Volume 1. SAS Institute Inc., Cary, N.C.

Wulder, M., S. Magnussen, P. Boudewyn and D. Seemann, 1999. Spectral variability related to forest inventory polygons stored within a GIS. In *IUFRO Conference on Remote Sensing and Forest Monitoring, Rogow, Poland, June 1 to 3, 1999* (in press).

Wynne, R.H., R.G. Oderwald, G. A. Reams and J.A. Scrivani, 2000. Optical Remote Sensing for Forest Area Estimation. *Journal of Forestry*, 98(5):31-36.