The Joint Remote Sensing Research Program
Publication Series

An Index of Annual Forest Clearing Using Paired SPOT-5 HRG Imagery

Author(s): Neil Flood[1,3], Tim Danaher[1,2], Peter Scarth[1,3], Tony Gill[1,2]

1. Joint Remote Sensing Research Program, University of Queensland
2. NSW Office of Environment and Heritage
3. Qld Department of Environment and Science

Date: November 2015; re-edited January 2020

Abstract: The government of the Australian state of New South Wales (NSW) has a program to monitor the clearing of trees across the whole state, using remotely sensed imagery, originally based around Landsat. In order to extend this program to use the higher resolution SPOT-5 sensor, an index was fitted to a large set of training data, mapped over previous years, to detect clearing in an automated fashion from bi-temporal imagery, separated by approximately one year. A number of different forms for this index were tested, and the best performing one was selected. The resulting index is used to produce an initial classification image showing areas likely to be clearing, which is then manually edited to produce a high accuracy map of clearing over a given year. This initial classification has been found to greatly ease the process of manual classification. The initial classification has a number of probability levels, which have user's and producer's accuracy for the clearing class of around 90%, while the user's and producer's accuracy for the non-clearing class are in excess of 99%. The index has been trained using data in eastern Australia, but would probably be of use in other parts of Australia.

Contents

1 Introduction

Under Australian constitutional arrangements, land management is largely the responsibility of the state governments. As a part of this responsibility, some states maintain laws governing the clearing of trees. Since 1995 the state government of Queensland has run a program of using satellite imagery to monitor broadscale tree clearing under the *Vegetation Management Act 1999* (QLD) (Queensland Government, 2015), and the government of the state of New South Wales (NSW) now has a similar program, under the *Native Vegetation Act 2003* (NSW) (New South Wales Government, 2015).

Much of this monitoring, in both Queensland and NSW, has been carried out using imagery from the Landsat Thematic Mapper (TM) and Enhanced Thematic Mapper (ETM+) instruments. The methods used are described by Danaher, Scarth, et al. (2010) and Scarth et al. (2008), but in simple terms, an initial automated classification is applied to multi-spectral images from a pair of dates (usually separated by one year). This initial classification is then manually edited by visual interpretation of the imagery, and any other information available, to produce a final classification which maps clearing events occurring in that year. The results of this classification are used to inform a range of government activities, including monitoring compliance with vegetation management legislation, and so a high degree of accuracy is required.

The Landsat instruments both have a ground pixel size of 30 m. While this is very good for mapping large areas, in NSW much of the vegetation now being cleared is comprised of woodlands and grasslands with scattered trees, where it can be difficult to detect small areas of clearing using the Landsat imagery. For this reason, a program was initiated in NSW to monitor using imagery from the Satellite Pour l'Observation de la Terre 5 (SPOT-5) High Resolution Geometric (HRG) instrument, which has 4 spectral bands similar to 4 of the Landsat TM/ETM+ instruments (TM bands 2, 3, 4 and 5), and a ground resolution of 10 m. In addition, it has a pan-chromatic instrument with nominal 2.5 m resolution.

The state of NSW covers a land area of 800 642 km^2 (Geoscience Australia, 2014). Mapping over such large areas is greatly aided by as much automation as possible, even when, as in this case, the area still needs to be checked manually as well. A great deal of this sort of large area monitoring has historically been done using the Landsat sensors, and coarser resolution instruments such as the Advanced Very High Resolution Radiometer (AVHRR) (Hansen et al., 2012; Coppin et al., 2004). Considerably less work on large area mapping has been done using the SPOT sensor.

Chirici et al. (2011) discuss the use of SPOT-5 imagery in Italy for mapping a very specific type of clearing in managed woodlands, known as coppice clearcutting. In this situation, dense managed forest is cut back to near ground level in small patches, creating fairly sharp contrast between clearing and surrounding forest, although because stumps are left in place, the trees quickly re-shoot, so monitoring must be done in a timely fashion. They compared multi-date image segmentation approaches and pixel-based change detection, using generic statistical methods, and found the object-based method to be of most value in delineating these sharp contrasts. Souza et al. (2003) have used the older SPOT-4 sensor to map forest degradation in the Amazonian rainforests of Brasil, using a spectral unmixing approach. We believe that the Australian savannah presents a unique set of challenges for automated methods, due to the great range of canopy densities to be found across the landscape.

For the first two annual periods in which the NSW Office of Environment and Heritage (OEH) carried out this mapping using SPOT-5 imagery, a simple clearing index was used for the initial classification, based on a simple radiometric standardisation procedure (Danaher, Hicks, et al., 2010). Subsequently, a more sophisticated radiometric standardisation procedure has been adopted (Flood et al., 2013), and combining this with the availability of the mapped clearing for those two periods to use as training data, it was felt appropriate to follow the methods previously used for Landsat by Scarth et al. (2008) and try to fit a more robust clearing index, to improve the initial classification.

This paper describes the methods used to create an automated classification of tree clearing for SPOT-5 HRG imagery, following the methods outlined by Scarth et al. (2008) for Landsat. This method has been used for three subsequent annual clearing periods, and the classifications have been manually edited, as discussed above. The resulting edited classifications have been used to assess the accuracy of the initial automated classification.

2 Methods

2.1 Overview

Scarth et al. (2008), working with Landsat TM/ETM+ imagery, used a large dataset of previously mapped clearing as training data to create a number of metrics which helped to predict clearing. We have tried to follow a similar approach, as far as possible. In the description below, we have included mention of all the components of that approach, and detailed which were found to be applicable to the available SPOT data, and which were not, and why they were excluded from the current work.

In general terms, these methods work in continuous data spaces, as opposed to working in terms of discrete classification. As noted by the reviews of Coppin et al. (2004) and Hansen et al. (2012), performing change detection with discrete classified images tends to have a higher error rate than working with continuous transformations, as the classification errors are compounded.

The first and most important part of the analysis by Scarth et al. (2008) was a function which takes the input reflectance in all spectral bands, for two dates, and calculates a value which is large for pixels which have probably been cleared, and small for pixels which have probably not been cleared. This index can then be thresholded at different levels to indicate increased likelihood of clearing, and manually interpreted in conjunction with any other information. This approach has been reproduced here. The imagery used is taken from the dry season for the region in question, in order to maximise the spectral difference between tree cover and grass cover. In the Australian context, most grasslands will brown off considerably during the dry season, in contrast to the trees. As discussed in the review by Coppin et al. (2004), the selection of time of year of imagery for change detection is dependant on the vegetation community and climatic regime being studied.

The second component of the Landsat methods of Scarth et al. (2008) was a similar function based around a model of overstorey foliage projective cover (FPC). Overstorey FPC is a measure of the tree cover, defined as "the vertically projected percentage cover of photosynthetic foliage from tree and shrub lifeforms greater than 2 m height", and has been modelled using Landsat multi-spectral reflectance imagery Armston et al., 2009. While no similar model has been fitted directly to SPOT data using coincident field survey data, Fisher et al. (2016) have used a cross-calibration approach to fit a model to the Landsat FPC model, i.e. using the Landsat predicted FPC in place of field data to create a model which predicts FPC from SPOT reflectance values. This is obviously less accurate than fitting directly to field data, but such field data was not available for the relevant dates and locations corresponding to available SPOT imagery. Some experiments with using this SPOT FPC model in the same way as Scarth et al. (2008) found that it gave no improvement over the simple spectral clearing index, and the form of that model is so similar to the spectral clearing index forms tested for the first component that this is not surprising. Consequently, in the current work, this FPC model has been used only indirectly, as a means of filtering the training data, as described in section 3.2.

The third component of the model of Scarth et al. (2008) used a long timeseries of FPC imagery from dry season dates, and further quantifies change in the most recent period relative to the variability of FPC over the preceding time series. This would fall into the class of "temporal trajectory" methods (Coppin et al., 2004). Since no such long time series of SPOT imagery was available, this component was not attempted for the current analysis.

2.2 Clearing Index Model

A range of forms were tested for a clearing index, as shown in equation 1.

$$\mathrm{CI} = a_0 + \sum_{k=1}^{2}\sum_{i=1}^{N} a_{i,k}\rho_{i,k} \tag{1a}$$

$$\mathrm{CI} = a_0 + \sum_{k=1}^{2}\sum_{i=1}^{N} a_{i,k}\ln(100\rho_{i,k}+1) \tag{1b}$$

$$\mathrm{CI} = a_0 + \sum_{k=1}^{2}\left[\sum_{i=1}^{N} a_{i,k}\ln(100\rho_{i,k}+1) + \sum_{i=1}^{N}\sum_{j=i}^{N} b_{i,j,k}\ln(100\rho_{i,k}+1)\ln(100\rho_{j,k}+1)\right] \tag{1c}$$

$$\mathrm{CI} = a_0 + \sum_{k=1}^{2}\sum_{i=1}^{N} a_{i,k}\rho_{i,k}^2 \tag{1d}$$

where CI is the clearing index, $N = 4$ is the number of reflective spectral bands in the SPOT-5 HRG instrument, $\rho_{i,k}$ is the surface reflectance of band i on the start/end date ($k = 1, 2$), and a_0, $a_{i,k}$ and $b_{i,j,k}$ are fitted coefficients. The use of the reflectance transformation $(100\rho + 1)$ in the logarithmic forms is primarily to avoid taking logarithm of zero. While it is possible that some optimisation of this offset may produce a better performing model, some experimentation by the authors suggests that it is not an important parameter.

For each pixel in the training data, a pixel which was manually classified as clearing was assigned $\mathrm{CI} = 1000$, while pixels not classified as clearing were assigned $\mathrm{CI} = 0$.

Following Scarth et al. (2008), the models were fitted to the training data using ordinary least squares regression, solved using singular value decomposition where singular values smaller than 0.001% of the largest singular value are discarded (as implemented in the SciPy routine `scipy.linalg.lstsq` (Jones et al., 2001)).

This procedure of fitting a linear regression to relate classified outputs to a space of continuous inputs is exactly equivalent to using Linear Discriminant Analysis to fit a two-class classifier (Hastie et al., 2009) (page 109).

The selection of the training data is described in section 3.2.

2.3 Evaluating Model Prediction

These models were fitted using the training data from the periods 2008–2009 and 2009–2010. The model which best fitted the training data was equation 1c, and this function was used in subsequent years to create an initial classification of clearing, which was then manually edited to create a final classification appropriate for the legislated requirements of the NSW government. This process has now been carried out over several years, and manually edited and checked clearing classification images are now available for the periods 2010–2011, 2011–2012 and 2012–2013 (Office of Environment and Heritage, 2014). This forms a more comprehensive dataset which can be used as independant data for validation. All of the model forms were tested with this validation data, and the results shown in section 4.

The performance of each classification model was evaluated using receiver-operator characteristic (ROC) curves. The ROC curve plots the performance of a classification index in the 2-dimensional space of its true-positive rate and its false-positive rate. Fawcett (2006) provides a good introduction to ROC analysis. In very simple terms, a continuous classification index is plotted as a curve in this space by varying the threshold at which we distinguish between a positive and a negative, i.e. clearing or non-clearing. Generally, the larger the area under this curve, the better the performance of the classifier.

Because the result of this classifier is still intended to be manually edited to create a final classification, little emphasis has been placed on selecting a single optimal threshold of the clearing index. Instead, the index is simply coded to a range of codes representing different levels of likelihood of clearing, and these codes are the starting point for editing.

This means that the commonly used measures of classification accuracy, such as user's and producer's accuracy, are more difficult to interpret in this context. In order to give some idea of the accuracy of the clearing index, the thresholds used to code the initial classification image are shown along with the corresponding user's and producer's accuracy measures. However, it is important to remember that these do not represent the final accuracy of clearing classification, but rather an indicator of the starting point for manual editing.

3 Data

3.1 Reflectance imagery

The SPOT-5 panchromatic and multispectral images were purchased as level 1A processed images and orthorectified by a range of commercial providers, using a SPOT-5 orbital model. These images were rectified to existing SPOT-5 orthorectified base images that were produced from 2004–2005 dates. The orthorectified base was produced by Geoimage Pty Ltd (`http://www.geoimage.com.au`), using ground control points from GPS, orthorectified photography and topographic data, and DEM's based on 1:25,000 topographic data where available, with Shuttle Radar Topography Mission (SRTM) data used in the western third of the state. The absolute accuracy was checked by analysts within OEH, and the root mean square error on a selection of 1200 well defined control points was found to be 1.5 m for the panchromatic band and 4.6 m for the multi-spectral bands. Image-to-image registration between dates was also visually checked by OEH analysts for each date pair, and was generally found to be well within one pixel.

The reflectance imagery was corrected to standardised surface reflectance, with a nadir view angle and a solar incidence angle of 45°. It has been calculated from the at-sensor radiance by correcting for atmospheric effects, the effects of the Bi-directional Reflectance Distribution Function (BRDF), and topographic illumination effects, using the methods described by Flood et al. (2013). This level of standardisation removes much of the variability not related to change of the land surface, meaning that it is more feasible to fit a single model which will detect clearing in a range of locations, with a variety of date pairs. In general, the benefits of this kind of image standardisation for change detection are discussed in more detail by Coppin et al. (2004) and Hansen et al. (2012).

Cloud and cloud shadow masks have been prepared using the methods outlined by Fisher (2014), and manually edited to remove all detected errors.

Topographic shadowing was masked using the methods described by Robertson (1989), using the topography given by the Shuttle Radar Topography Mission (SRTM) at 30 m resolution (Farr et al., 2007; Gallant et al., 2009). Pixels with solar incidence angle greater than 80° to the surface normal were also masked as being affected by topographic shadowing.

The surface reflectance imagery was sharpened to 5 m pixel size using the panchromatic band. The panchromatic imagery was degraded (by averaging) from its nominal 2.5 m pixel size to 5 m, as this was believed to be the level at which it was accurately co-registered to the multi-spectral imagery. This 5 m panchromatic imagery was used to sharpen the multi-spectral surface reflectance imagery, using a simple in-house method designed to preserve the radiometric integrity of the reflectance values. This method uses a Theil-Sen estimator (Sen, 1968; Wilcox, 2010) on a local window to predict the higher resolution value from the lower resolution and the panchromatic value. It has not yet been formally described elsewhere. The clearing index has been tested both with and without the pan-sharpening, with similar results, and this aspect is not considered crucial to the method. It is mentioned here merely for completeness.

3.2 Training Data

For the periods 2008–2009 and 2009–2010, the NSW OEH have mapped forest clearing by applying preliminary change methods to pairs of SPOT-5 HRG imagery (Danaher, Hicks, et al., 2010; Office of Environment and Heritage, 2012). The resulting mapped clearing was used as training data for the current work.

Briefly, the original mapping began with a classification based on a simple differencing method (Danaher, Hicks, et al., 2010). This was then manually edited by visual inspection of both dates, and any other ancillary imagery available. This resulted in raster images with a number of different clearing classes. All classes denoted clearing or non-clearing, and different codes indicated the operator's interpretation of the intended land use after clearing, such as Urban, Cropping, etc. All clearing codes represented removal of established woody vegetation (i.e. trees) (Office of Environment and Heritage, 2012). It should be emphasised that there are no codes in the edited imagery which have not been manually recoded in some fashion. In other words, an operator has visually inspected every part of the image, so that there is some reason to believe that the accuracy of the final classification is very high, although this has not been quantified.

All image pairs in which more than 0.1% of the area was classified as clearing were selected as

contributing training pixels. This percentage is typical, as in any given annual period, only a very small fraction of the area has been cleared (Office of Environment and Heritage, 2012). In order to cope with the fact that only a very low proportion of pixels represent clearing, the selection of pixels handled clearing and non-clearing pixels slightly differently. Two regular grids were layed onto the imagery. One grid had a spacing of 500 m, the other a spacing of 100 m. From the 500 m grid, only non-clearing pixels were selected, while from the 100 m grid, only clearing pixls were selected. This means that clearing pixels were over-sampled, relative to non-clearing pixels, in order to better represent the range of vegetation being cleared.

Points for which the FPC on the start date of the pair was less than 8% were excluded on the basis that these are almost certainly not clearing, since such a low value of FPC suggests that there are no trees present to begin with. Since the Australian landscape has large areas of grassland with very few trees, there are very large numbers of such pixels. Inclusion of these pixels in the training data has the potential to confound the fitting of the clearing index because the process will be dominated by the correct, but uninteresting, predictions that areas with no trees were not cleared. The same exclusion was used successfully by Scarth et al. (2008).

This created a set of 1763465 pixels, of which 4774 were clearing and 1758691 were not clearing. Note that without the oversampling of clearing pixels, there would be only approximately 200 clearing pixels in the sample, which was felt to be too small to be representative. However, it is important to preserve the general fact that most pixels are not cleared, in order to weight heavily against errors of commission. Because of the large areas in the landscape which are not cleared in any given period, even a small false positive rate (i.e. classifying non-clearing pixels as clearing) quickly results in very large areas being mis-classified. This does mean that the total set of data was very large, but it remained manageable.

3.3 Validation Data

A validation dataset was extracted by a similar procedure, but the pixel extraction grid was offset from that used for the training data, so these pixels are independent of the training data. The validation data was taken from the much larger dataset, including annual clearing from the 5 annual periods between 2008 and 2013. Images with almost no clearing were omitted, as with the training data, but the clearing was not oversampled relative to the non-clearing pixels, nor were low-FPC pixels excluded. The pixel selection grid had a spacing of 500 m, and was offset from the grid used for training pixels by 250 m in both X and Y directions. Thus this validation dataset is fairly representative of the whole landscape, not just the part used for fitting the model. The total number of pixels in this validation dataset is 5808918, with 4271 clearing pixels.

This validation set of pixels is used for all accuracy assessment statistics shown in what follows. As mentioned in section 3.2, the classification used as ground truth has been visually inspected in great detail, and is assumed to be of very high accuracy.

4 Results

The different models given in equation 1 were evaluated against the validation dataset, by calculating ROC curves. For each model, the area under the ROC curve is given in table 1, along with the abbreviation used for each model. These abbreviations identify the individual ROC curves, plotted in figure 1.

Equation	Abbreviation	Description	ROC Area
1a	bn	Reflectance in band n, for all n	0.9917
1b	lbn	Log(reflectance), band n, for all n	0.9943
1c	lbn,lbnm	Log(reflectance), band n, and products of all pairs	0.9963
1d	bnsq	Square of reflectance in band n, for all n	0.9785

Table 1: Areas under ROC curves, for each clearing model tested. Abbreviations are used in figure 1 to identify each curve

The ROC curves shown in figure 1, along with the area calculations in table 1, both suggest that the best performing model is the one selected during the original fitting process, i.e. equation 1c (see section 2.3). The final equation, including the fitted coefficients, is given in equation 2.

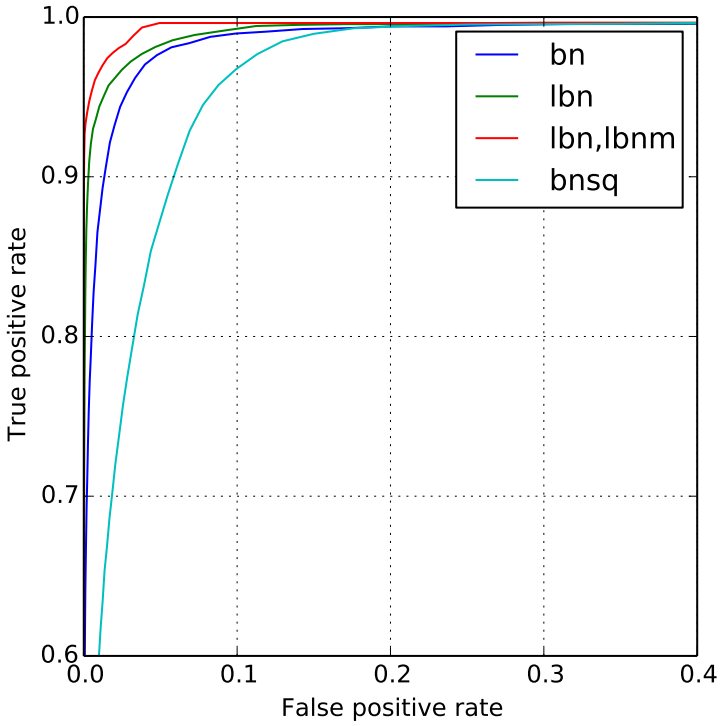

Figure 1: Receiver-Operator Characteristic Curves for the different clearing index models tested. The model labels in the legend are the abbreviations found in table 1, linking them to the relevant equation numbers. Note that only the top-left part of the plot is shown.

$$
\begin{aligned}
CI \quad = \quad & 6.1477892 + 14.7004397R_{1,1} - 85.6395164R_{2,1} + 79.1790298R_{3,1} + 20.8942184R_{4,1} - \\
& 28.0708718R_{1,2} + 99.6591326R_{2,2} - 112.3720233R_{3,2} + 13.3256975R_{4,2} - \\
& 20.7211726R_{1,1}R_{1,1} + 71.3091213R_{1,1}R_{2,1} - 3.4108285R_{1,1}R_{3,1} - 17.6360425R_{1,1}R_{4,1} - \\
& 54.9295183R_{2,1}R_{2,1} + 19.2063403R_{2,1}R_{3,1} + 32.1737616R_{2,1}R_{4,1} - 11.5581789R_{3,1}R_{3,1} - \\
& 12.6196460R_{3,1}R_{4,1} - 12.2235715R_{4,1}R_{4,1} + \\
& 22.7935147R_{1,2}R_{1,2} - 76.3814644R_{1,2}R_{2,2} + 19.9753522R_{1,2}R_{3,2} + 2.2928294R_{1,2}R_{4,2} + \\
& 46.0685420R_{2,2}R_{2,2} - 27.7205495R_{2,2}R_{3,2} - 10.5163884R_{2,2}R_{4,2} + 16.1280469R_{3,2}R_{3,2} + \\
& 10.4288073R_{3,2}R_{4,2} + 4.6311733R_{4,2}R_{4,2} \qquad (2)
\end{aligned}
$$

where $R_{i,k} = \ln(100\rho_{i,k} + 1)$ and $\rho_{i,k}$ is surface reflectance for band i on date k, k is 1 for the start date and 2 for the end date.

An example of how this index looks on an image is shown in figure 2. The top pair of images ((a) and (b)) are the multi-spectral reflectance imagery for the start and end dates. The bottom left image (c) is the clearing index calculated for this date pair, and the bottom right image (d) is the coded values resulting from applying a range of thresholds to the clearing index, to indicate different likelihoods of a given pixel having been cleared. In the reflectance imagery, the band combination is (red, green, blue) = (bands 4, 3, 1). With these bands, dense forest shows as dark green, green grass shows as pale green, dry grass shows as purple, with bare ground showing darker purple.

The thresholds used to code the clearing index, as shown in figure 2(d), are somewhat arbitrary, but were chosen to show increasing levels of likelihood that a given pixel has been cleared. As mentioned earlier, the intention of this coded image is as a starting point for a manual editing process. Manual editing is performed on this categorical image, rather than on the continuous index, by simply recoding categorical values within visually outlined areas of interest. Obviously this is easier on screen, when the images can be overlayed on top of each other. Thresholds were chosen to slightly overestimate vegetation change so no clearing is missed but not to grossly overestimate change which would make the visual editing too onerous for the operator. Given that the number of pixels which are actually

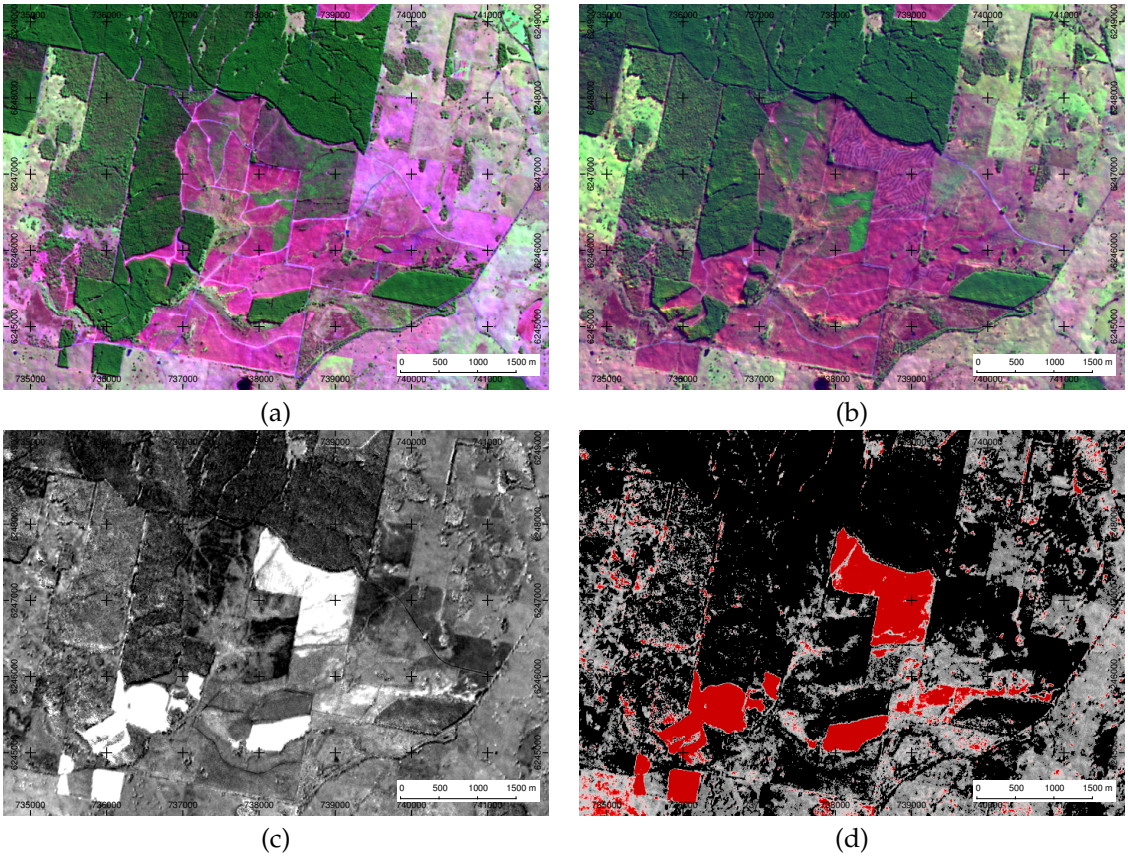

Figure 2: Example imagery showing the clearing index in action. The top two images are the multi-spectral reflectance (red,green,blue=bands 4,3,1). Dark green is dense forest, light green is grassland, purple is dry and patchy grass. (a) Reflectance on start date; (b) Reflectance on end date; (c) Clearing index (grayscale, with white being most likely clearing); (d) Coded clearing index. Red areas are most likely clearing, and gray levels from light to dark are progressively less likely. Black is "not clearing".

cleared is very small compared to the number of pixels which are not cleared (on the order of 0.1%), even a very small false positive rate results in a fairly large number of false positive pixels, which must then be manually edited out again. Table 2 shows the thresholds used, and the corresponding true positive and false positive rates, an the number of false clearing pixels, based on the validation data set.

To present equivalent information in a slightly different format, perhaps more familiar in the context of accuracy assessment, the user's and producer's accuracies are shown for each of these thresholds, in table 3. Note that the clearing producers accuracy is equal to the true positive rate, and the false positive rate is equal to $(100-$non-clearing producer's accuracy$)$. The accuracies for the non-clearing class are inevitably very high, as most pixels are not cleared, and this is easily predicted by the classification.

4.1 The Value of Band 4

In addition to the analyses conducted above, one additional variant of equation 1c was also analysed briefly. The work described in this paper is all based on the SPOT-5 HRG instrument, which has 4 reflective wavelength bands. The equivalent instrument on the successor platforms SPOT-6 and SPOT-7 has one extra band, in the blue wavelengths, but does not have an equivalent of the HRG band 4, in the shortwave infra-red (SWIR). While we cannot test directly what will be possible with this new instrument until we have sufficient imagery coincident with mapped clearing, we can conduct a simple test on the impact of removing the SWIR band. We re-fitted the selected clearing index model (equation 1c) but using only the first three bands, leaving out the data from band 4 (SWIR). As with the other models tested, this was fitted using the same training data and tested against the same validation data. The

Threshold	True Positive (%)	False Positive (%)	False Clearing Pixels
14.28	97.799	1.913	111015
18.28	95.903	0.651	37803
22.28	93.070	0.069	3981
26.28	86.654	0.006	348
29.28	81.035	0.004	245
31.78	76.165	0.004	236
33.78	70.733	0.004	235
36.28	63.779	0.004	235

Table 2: Thresholds used to code clearing index. The rates of true and false positive are calculated from the thresholds against the validation data set. The number of false clearing pixels can be compared with the number of true clearing pixels in the validation data set, which is 4127.

Threshold	Clr User's	Clr Producer's	Non-clr User's	Non-clr Producer's
14.28	3.626	97.799	99.998	98.087
18.28	9.776	95.903	99.997	99.349
22.28	49.962	93.070	99.995	99.931
26.28	91.405	86.654	99.990	99.994
29.28	93.389	81.035	99.986	99.996
31.78	93.236	76.165	99.982	99.996
33.78	92.783	70.733	99.978	99.996
36.28	92.058	63.779	99.973	99.996

Table 3: User's and producer's accuracy for clearing and non-clearing classes, for a range of thresholds of the clearing index, calculated against the validation data set. Accuracies are expressed as percentages.

resulting ROC curve is shown in figure 3. Also included in this figure is the ROC curve for the same model using all bands, as it appeared in figure 1, for comparison.

The removal of the SWIR band causes a substantial change in the accuracy of the resulting clearing index. The curve pulls away from the left hand side much earlier, indicating a much greater false positive rate for a given true positive rate. This would make manual editing of the resulting classified image much more difficult.

5 Discussion

The fitted clearing index (equation 2) is able to reliably detect around 90% of the clearing which occurs between two imagery dates, as shown by the ROC curves in figure 1, with selected details in table 2. The main trade-off arises when we try to threshold the index so as to detect the remaining clearing, as this is when the false positive rate starts to rise to unacceptable levels.

The details in table 2 highlight the problems associated with even a very small false positive rate, which is due to the fact that the overwhelming majority of pixels are not cleared in any given annual period. Because of this, anything other than a very nearly perfect classifier will still give rise to unmanageably large numbers of false positives. This suggests that there are very real practical limits on the ability of any similar classifier to be reliable.

With that in mind, the actual performance of the clearing index across the whole landscape, as represented by the validation data set, is very encouraging. As discussed in earlier sections, it is not expected or intended that this clearing index can be used to create a final classification to map clearing, but rather to produce an initial set of classes which can be manually recoded by visual interpretation. The index of equation 2 has been in use for this purpose for the three clearing periods since it was fitted (2010–2011, 2011–2012 and 2012–2013), and has been found to be well suited to this mode of operation.

All of the clearing index models tested show some ability to predict clearing. However, as visible on the ROC curves, some of the models pull away from the left hand vertical axis much earlier than others. This corresponds to an increase in the false positive rate at a much lower true positive rate. For example,

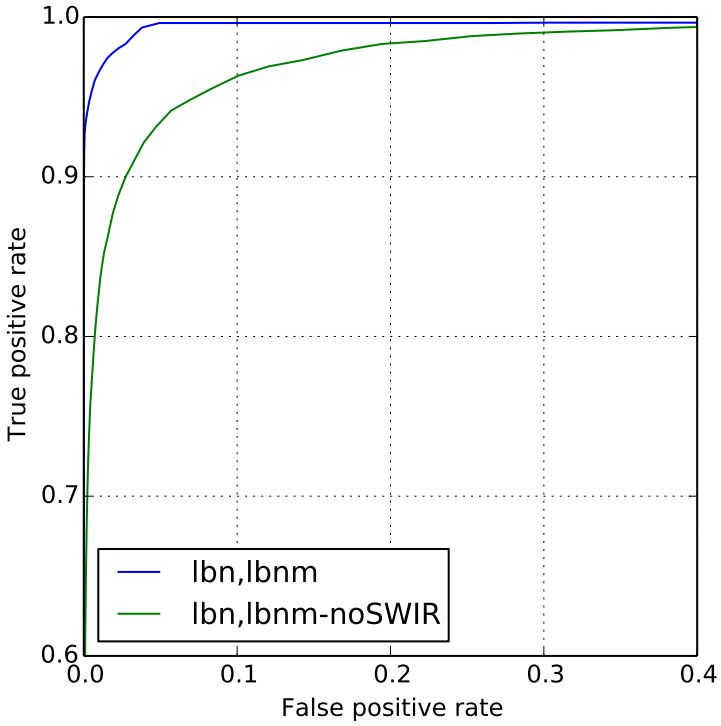

Figure 3: Receiver-Operator Characteristic Curves for the clearing index using equation 1c, with and without including band 4 (SWIR).

the "bnsq" model (equation 1d) shows a false positive rate of around 1% for a true positive rate of only 60%, whereas the selected model does not reach this level of false positives until the true positive rate is over 95% (see table 2). The ROC curve for the selected model hugs the left hand vertical axis much more tightly, and it is this characteristic which makes it the most useful for the intended purpose.

It is worth noting that the final form of the index function is very similar to the form found for the FPC index (which is a separate model). This reflects the fact that both are intended to detect tree cover, although in the case of FPC it is to measure the amount of cover on a single date, while the clearing index is intended to detect removal of tree cover between two dates, and so has coefficients tuned for this precise purpose. However, the fact that both models involve a logarithmic transformation of reflectance data, and both are quadratic in the transformed reflectance, suggests that these transformations are likely to be of value in other work relating reflectance to vegetation cover.

The example imagery in figure 2 shows a number of features of interest. In the south-west quadrant are several patches of clearing of dense forest. These show as the brightest areas in the clearing index, figure 2(c), and are coded red in figure 2(d). Scattered across the image are small patches which show as brighter in figure 2(c), and some which are bright enough to be coded as red in figure 2(d), but which are probably not clearing, as seen by comparing the two reflectance images. These represent false positives which would need to be eliminated in the manual editing process. Just north of the centre of the example area is a large patch which is coded as red, for clearing, and appears on the reflectance imagery as removal of some significant vegetation. The marks of agricultural machinery are clearly visible in the end-date reflectance in figure 2(b). However, the vegetation shown in the start-date reflectance is not consistent with dense forest, and represents less dense vegetation such as low shrubs and young trees. Whether this should be classified as clearing must be decided during the manual editing process.

The index has been fitted for the vegetation communities of the state of New South Wales, in eastern Australia. Similar vegetation communities exists in many parts of Australia, and there is good reason to believe that the spectral characteristics are sufficiently similar that the model fitted here would work equally well in other parts of the country. For example, the Landsat FPC index fitted by Armston et al. (2009) was originally fitted on field data from the state of Queensland, further north, but has been found (anecdotally) to work equally well in most of eastern Australia. The fractional vegetation cover model of Scarth et al. (2010) was fitted using a single set of field data from across much of the country, and found to apply well in most areas. These experiences suggest that the spectral characteristics which dis-

tinguish woody vegetation are sufficiently similar across much of the country. Whether this is true on other continents is not immediately clear, and some testing of this assumption is recommended before applying it in other parts of the world. In addition, the woody vegetation of Australia is dominated by evergreen species, and so mostly not subject to the limitation of seeing deciduous tree species during leaf-off conditions. It seems certain that end-date observations of trees in leaf-off condition would produce false positives for clearing.

On a related note, in areas in which grasses are less subject to browning off during dry periods, it is possible that the ability of this index to spectrally distinguish tree cover from herbaceous cover will be reduced. This would have the twin effects of making it: (a) harder to detect clearing when it happens, giving reduced true positives; and (b) easier to mistake a reduction in grass cover for clearing of trees, producing increased false positives. Anecdotally, this type of behaviour has been observed when using the index with either the start or end image from a season which was wetter than normal.

No testing has been done of the sensitivity of the fitted model to the precise detail of the radiometric processing applied to the input imagery. To use the model as fitted would almost certainly require an atmospheric correction, due to the significant shifts in reflectance values resulting from removal of the atmospheric signal. Because of such offsets, it is not clear that a simpler, relative correction, such as Dark Object Subtraction, would be sufficient, as the baseline image would still contain the effects of the path radiance, and so the model coefficients presented here would not be appropriate.

In order to use at-sensor (top-of-atmosphere) reflectance as an input, without any atmospheric correction, it seems likely that: (a) new coefficients would need to be fitted; and (b) the result would be less reliable, due to the increased variability of the input reflectance values, due to variations in atmospheric path length and similar effects.

The work described here is aimed at using the SPOT-5 HRG sensor. However with this platform reaching its end of life in 2015, and the replacement platforms SPOT-6 and SPOT-7 having a slightly different instrument, it was felt timely to take the opportunity to assess how much this work might be affected by the change in instrument, and whether it will be a viable option for the future of this work. The simple test results described in section 4.1 suggest that the loss of band 4 will have serious impacts on this work. Until we have access to sufficient imagery from the new sensors, coincident with mapped clearing, we are not in a position to make a definitive statement about this, but the degradation of the clearing index shown in figure 3 does suggest that other alternative sensors should be explored as well.

6 Conclusions

An existing dataset of mapped tree clearing has been used to create an index which indicates tree clearing from a pair of annual SPOT-5 images. Depending on what thresholds are then applied to this index, it has been shown to map around 93% of the tree clearing, while having a false positive rate of less than 0.07% (table 2). It has been used across broad areas of the New South Wales landscape, providing an automated initial classification of clearing, for use as a starting point for a high-accuracy, manually edited map of tree clearing. This process has been carried out for three successive one-year intervals. While the very low false positive rate is highly desirable, this should be viewed with some caution, as the relative rarity of clearing within the landscape means that even quite low false positive rates can still imply a large number of pixels falsely classified as clearing, and thus requiring manual editing. However, with this in mind, carefully chosen thresholds on the clearing index have been found to create a very useful initial classification.

The dependence of this index on the presence of the short wave infra-red band of SPOT-5 has been demonstrated, suggesting the need for caution in attempting to use a similar approach with the newer SPOT-6/7 sensors which do not include an equivalent band.

Acknowledgments

This work was supported by the Joint Remote Sensing Research Program, a collaboration of the state governments of Queensland, New South Wales and Victoria. The edited clearing images used for training and validation were created by a dedicated team of operators within the NSW Office of Environment and Heritage, and this work would not have been possible without them.

References

Armston, John D., Robert J. Denham, Tim J. Danaher, Peter F. Scarth, and Trevor N. Moffiet (2009). "Prediction and validation of foliage projective cover from Landsat-5 TM and Landsat-7 ETM+ imagery". In: *Journal of Applied Remote Sensing* 3.1, pp. 033540–033567. DOI: 10.1117/1.3216031. URL: http://dx.doi.org/10.1117/1.3216031.

Chirici, Gherardo, Diego Giuliarelli, Daniele Biscontini, Daniela Tonti, Walter Mattioli, Marco Marchetti, and Piermaria Corona (2011). "Large-scale monitoring of coppice forest clearcuts by multitemporal very high resolution satellite imagery. A case study from central Italy". In: *Remote Sensing of Environment* 115.4, pp. 1025–1033.

Coppin, P., I. Jonckheere, K. Nackaerts, and B. Muys (2004). "Digital change detection methods in ecosystem monitoring: a review". In: *International Journal of Remote Sensing* 25.9, pp. 1565–1596.

Danaher, Tim, Richard Hicks, Tony Gill, Geoff Horn, Arndt Meier, Stuart Smith, and Andy Taylor (2010). "High Resolution Monitoring of Woody Vegetation Change over New South Wales". In: *Proceedings of the 15th Australasian Remote Sensing and Photogrammetry Conference, 13–17 September, 2010, Alice Springs, Australia*, pp. 225–236.

Danaher, Tim, Peter Scarth, John Armston, Lisa Collett, Joanna Kitchen, and Sam Gillingham (2010). "Remote Sensing of Tree-Grass Systems: The Eastern Australian Woodlands". In: *Ecosystem Function in Savannas: Measurement and Modelling at Landscape to Global Scales*. Ed. by Michael J. Hill and Niall P. Hanan. CRC Press, Boca Raton, Florida, USA, pp. 175–193.

Farr, Tom G. et al. (2007). "The Shuttle Radar Topography Mission". In: *Reviews of Geophysics* 45.2, RG2004.

Fawcett, Tom (2006). "An introduction to ROC analysis". In: *Pattern Recognition Letters* 27.8, pp. 861–874.

Fisher, Adrian (2014). "Cloud and cloud-shadow detection in SPOT5 HRG imagery with automated morphological feature extraction". In: *Remote Sensing* 6.1, pp. 776–800.

Fisher, Adrian, Michael Day, Tony Gill, Adam Roff, Tim Danaher, and Neil Flood (2016). "Large-area, high-resolution tree cover mapping with multi-temporal SPOT5 imagery, New South Wales, Australia". In: *Remote Sensing* 8.6, p. 515.

Flood, Neil, Tim Danaher, Tony Gill, and Sam Gillingham (2013). "An Operational Scheme for Deriving Standardised Surface Reflectance from Landsat TM/ETM+ and SPOT HRG Imagery for Eastern Australia". In: *Remote Sensing* 5.1, pp. 83–109.

Gallant, J.C. and A. Read (Aug. 2009). "Enhancing the SRTM data for Australia". In: *Proceedings of Geomorphometry*.

Geoscience Australia (2014). *Area of Australia - States and Territories*. URL: http://www.ga.gov.au/scientific-topics/geographic-information/dimensions/area-of-australia-states-and-territories (visited on Dec. 2014).

Hansen, Matthew C. and Thomas R. Loveland (2012). "A review of large area monitoring of land cover change using Landsat data". In: *Remote Sensing of Environment* 122, pp. 66–74.

Hastie, Trevor, Robert Tibshirani, and Jerome Friedman (2009). *Elements of Statistical Learning: Data Mining, Inference, and Prediction*. 2nd. Springer.

Jones, Eric, Travis Oliphant, Pearu Peterson, et al. (2001). *SciPy: Open source scientific tools for Python*. URL: http://www.scipy.org/ (visited on May 22, 2015).

New South Wales Government (2015). *Native vegetation reports and resources*. URL: https://www.environment.nsw.gov.au/topics/animals-and-plants/native-vegetation/reports-and-resources/reports (visited on Jan. 2020).

Office of Environment and Heritage (Feb. 2012). *NSW Annual Report on Native Vegetation 2010*. Tech. rep. OEH 2011/0685. 59–61 Goulburn Street, Sydney PO Box A290, Sydney South, NSW 1232: New South Wales Office of Environment and Heritage. URL: http://www.environment.nsw.gov.au/resources/vegetation/110685nvar2010.pdf.

— (Mar. 2014). *NSW Report on Native Vegetation 2011–13*. Tech. rep. OEH 2014/0270. 59–61 Goulburn Street, Sydney PO Box A290, Sydney South, NSW 1232: New South Wales Office of Environment and Heritage. URL: http://www.environment.nsw.gov.au/resources/vegetation/2011-13NSWAnnRepNatVegFinal.pdf.

Queensland Government (2015). *Statewide Landcover and Tree Study (SLATS)*. URL: https://www.qld.gov.au/environment/land/vegetation/mapping/slats (visited on Nov. 2015).

Robertson, P.K. (1989). "Spatial transformations for rapid scan-line surface shadowing". In: *Computer Graphics and Applications, IEEE* 9.2, pp. 30–38. ISSN: 0272-1716. DOI: 10.1109/38.19049.

Scarth, Peter, Sam Gillingham, and Jasmine Muir (Sept. 2008). "Assimilation of spectral information and temporal history into a statewide woody cover change classification". In: *14th Australasian Remote Sensing and Photogrammetry Conference*. URL: http://dx.doi.org/10.6084/m9.figshare.96338.

Scarth, Peter, Achim Röder, and Michael Schmidt (2010). "Tracking Grazing Pressure and Climate Interaction - The Role of Landsat Fractional Cover in Time Series Analysis". In: *Proceedings of the 15th Australasian Remote Sensing and Photogrammetry Conference, 13–17 September, 2010, Alice Springs, Australia*, pp. 936–948.

Sen, Pranab Kumar (1968). "Estimates of the regression coefficient based on Kendall's tau". In: *Journal of the American Statistical Association* 63.324, pp. 1379–1389.

Souza, Carlos, Laurel Firestone, Luciano Moreira Silva, and Dar Roberts (2003). "Mapping forest degradation in the Eastern Amazon from SPOT 4 through spectral mixture models". In: *Remote Sensing of Environment* 87.4, pp. 494–506.

Wilcox, Rand R (2010). *Fundamentals of modern statistical methods: Substantially improving power and accuracy*. Springer Science & Business Media.