

Identifying Bare Soil Using Satellite Imagery and Supervised Image Classification.



Jim Walsh: 77677316

13/10/2020

Acknowledgements

I would like to thank Professor Rien Visser, Campbell Harvey, and Stephan Hoffmann for their guidance throughout this project. A special thanks goes to Mark Forward and Reihana Fisher for taking the time to assist with data collection out in the field. Finally, I would like to thank my fellow classmates for their support and input throughout the year.

Table of Contents

Acknowledgements.....	2
1. Executive Summary.....	4
2. Introduction	5
3. Literature review.....	6
3.1 Erosion and soil disturbance	6
3.2 Satellite imagery	7
3.3 Image processing	8
4. Problem statement and Objective.....	10
5. Methodology.....	10
5.1 Site locations	10
5.2 Ground-truth data.....	12
5.3 Analysing satellite images	13
5.4 Creating and evaluating training samples.....	14
5.5 Training sample sensitivity analysis	17
6. Results and Data analysis.....	21
7. Discussion.....	23
8. Conclusion.....	26
9. References	27

1. Executive Summary

Forestry soil disturbance assessments are generally carried out using ground-based methods such as plots and line transects, these methods are labour intensive and time consuming. With the increasing availability of higher resolution daily satellite imagery and improving image classification tools, there is the opportunity to use these resources to change the way site disturbance is measured.

Seven harvest sites in the South-Island of New Zealand were used to assess the accuracy of using satellite images and supervised image classification to identify soil disturbance of forestry sites. Three of the sites were around Nelson and the other four were around Christchurch. Ground truthing was carried out at each site, which involved creating two lines of 15 points each where the points were spaced at 10 meter intervals. The location of each point was recorded using a Trimble Geo7X handheld together with a Trimble Zephyr 3 GPS receiver, which provided an accuracy of 10-15 cm. At each point, the major land cover type was determined by visual inspection to be either bare soil or not bare soil (e.g. slash or vegetation). This data was recorded to later be compared against the classified image.

Satellite images were obtained through the PLANETSCOPE constellation with a 3m by 3m resolution, the images used were ideally taken closest to the date the ground-truthing data was collected for each site as well as being the clearest image with minimal shading and/or cloud cover. The images were uploaded into ArcMap Desktop 10.7.1 and processed using two supervised classification tools, Maximum Likelihood Equation (MLE) and Support Machine Vector (SVM). Once the images were classified the ground-truth points were uploaded and projected over the classified images. This allowed for each ground-truth point to be compared to the individual pixel it represented and for an accuracy assessment to be carried out for each site.

The accuracy assessment allowed for comparison between classification methods and techniques used within the classifications. The average overall agreement for MLE and SVM was 64% and 65% respectively, while the Kappa coefficient was 0.19 and 0.18, respectively. However, when looking at the best-case scenarios for each classification technique the average overall agreement for MLE and SVM was 68% and 72% respectively, while the Kappa coefficient was 0.23 and 0.27, respectively. Confirming that the SVM classifier generally outperforms the MLE classifier.

When using the SVM classifier it generally returned a higher User's accuracy compared to the MLE classifier, however, the MLE classifier provides a higher Producer's accuracy. This means that the maps created using the SVM classifier are more accurate from the map users' point of view. The supervised classification techniques were able to easily identify large disturbances such as roads and skid sites but struggled to pick up smaller disturbances due to the effects of 'mixed' pixel, where the pixels contain more than a single class.

When assessing the accuracy of the classification, ideally ground-truth data would be collected with a random scatter over the entire site, which would help provide a more significant accuracy assessment. As higher resolution imagery becomes more readily available from both daily satellite images or drones and fixed wing aircrafts, there is the opportunity to increase the accuracy of the classifications.

2. Introduction

The plantation forest industry covers approximately 1.69 million hectares, or 7% of New Zealand's total land area (Ministry for Primary Industries, 2019). For most of a plantation forest's life the land is covered with trees and it is well known that vegetation, and in particular trees, improve slope stability and reduce erosion (Norris et al., 2008). Tree roots reinforce soil making it stronger, and the tree canopy tends to make the soil drier (through hydrological processes of interception and transpiration), which also increases soil strength (Phillips et al., 2015). However, there is a time when this land cover is dramatically changed from the harvesting of the trees. Clear-fell harvesting effectively eliminates evapo-transpiration, often resulting in saturated soils and increased subsurface water flow, thus increasing the risk of slip erosion (Elliot et al. 1999). The loss of canopy cover exposes the soil to direct rainfall impact, increasing the amount of fluvial erosion. Post-harvest is when the land is most vulnerable to erosion and debris flows, which can have a range of negative impacts onsite and offsite.

Exposed soil from forestry practices such as earth works, harvesting, and mechanical land preparation have the potential to create sediment issues that do not meet the regulations outlined under the National Environmental Standards for Plantation Forestry (NES-PF). The NES-PF states that sediment originating from earthworks, harvesting, and mechanical land preparation must be managed to ensure that after reasonable mixing it does not give rise to any of the following effects on receiving waters: a) any conspicuous change in colour or visual clarity; b) the rendering of fresh water unsuitable for consumption by farm animals; c) any significant adverse effect on aquatic life (Resource Management (National Environmental Standards for Plantation Forestry) Regulations, 2018).

The amount and type of land cover is an important characteristic to measure when collecting soil disturbance data post-harvest. Land cover data can be used for tracking invasive species, plant growth, woody debris loadings, and the risk of surface water runoff and soil erosion (Kvamme, 2010). Creating land cover maps makes it easy to measure and compare factors across different study sites, for example, comparing different management practices such as erosion control measures and how the different practices can impact each site.

Determining land cover of forestry sites has historically been carried out using ground-based methods such as line intersects and plots, of which both methods are labour intensive and time consuming. These methods are still used today as they are tried and tested methods that are well known and provide consistent results. Advancements in technology have led to cheaper and more effective methods for collecting imagery, it has become a tool that is used daily by many industries, images of large areas can be taken from satellites and fixed wing aircrafts. Imagery processing software has also advanced and is used all around the world, software is readily available through the internet with a range of different processing tools.

This project will look at classifying the land cover of forestry sites post-harvest using satellite imagery. The images will be classified using the Maximum Likelihood classification and Support Machine Vector classification on ArcMap 10.7.1. Accuracy assessments of the classifications will be carried out using ground truth points from the harvest sites and comparing the ground-truth data to the classified images. The report will provide a critical review of the literature related to this project, identifying gaps in knowledge that can be filled. It will then provide a clear problem statement and objective of this project, provide a detailed methodology section on how the analysis will be carried out, a results and data analysis section, and a discussion and conclusion on the results obtained throughout the study.

3. Literature review

3.1 Erosion and soil disturbance

The current preferred harvesting practice in New Zealand is larger scale clear-cutting, based on logistical and economic benefits (Visser et al., 2018). The loss of the tree canopy results in a higher erosion risk due to more rainfall reaching the soil, less litter covering the soil surface, and the loss of root stabilisation.

There are two main types of erosion that occur within forestry operations, they are fluvial erosion and mass movement erosion (Satchell, 2018). Fluvial erosion is surface erosion caused by water scouring, it can be broken into the following categories. Sheet erosion, which is the removal of the thin layer of surface soil. Sheet erosion can become Rill erosion by scouring out small channels that are less than 60 cm deep and 30 cm wide. This Rill erosion can then become Gully erosion and Tunnel-gully erosion. Gully erosion is like Rill erosion except they are just channels larger than 60cm deep and 30cm wide. Tunnel-gully erosion forms tunnels under the ground surface that then collapse and become gullies.

Mass movement erosion is gravitational driven movement of soil mass downhill such as landslides and debris flows. They are generally a result of improper placement of slash or cutting of slopes for roads, which cause an increase in the stress placed on soil. The likelihood of mass movement erosion occurring is increased with heavy rainfall events that saturate the soil decreasing its strength.

Sheet erosion can be reduced with slash cover which stops the displacement of soil by rain drops. The smaller slash, such as needles and small branches, also provide nutrient recycling to the soil. However, larger slash, such as large branches and stem wood, can cause significant damage if they are not properly managed and become mobilised during a storm event.

Harvesting soil disturbances are caused by wheeled and tracked harvesting machinery, earthworks, and dragging logs across slopes (Satchell, 2018). Harvesting systems that create higher levels of soil disturbance produce higher levels of surface erosion. This surface erosion transports sediment downhill that either gets trapped by slash or enters waterways. Although the amount of sediment entering waterways from surface erosion is not as substantial when compared to mass movement, it occurs more often and can be brought on from smaller rainfall events than ones that trigger mass movement.

There are several rules under the NES-PF that regulate sediment entering waterways during different processes while harvesting. These rules are in place due to negative effects caused by increased sediment in waterways, such as conspicuous change in colour, rendering fresh water unsuitable for consumption by farm animals, and impacting aquatic life (Resource Management (National Environmental Standards for Plantation Forestry) Regulations, 2018). If overall soil disturbance can be reduced, so can the amount of sediment entering waterways.

A study carried out by Petherick (2014), looked at the amount of long-term unproductive land as a proportion of the total harvest area for forestry operations in the South Island of New Zealand. This unproductive area was classed as landings and permanent forestry roads (roads used for accessing skids, not skid trails). This was done for harvest sites with at least eight skid sites. Petherick found an average unproductive area of 4.82% of the total harvest area. He found this unproductive area using satellite images and ArcGIS to measure the area of the skids and lengths of roads within the harvest area, the roads' lengths were then multiplied by a width depending on the type of road to get a total sum of unproductive area for skids and roads.

A typical soil disturbance survey is completed using two sampling methods, Grid Point-Intercept (GPI) and the Traverse method (Curran et al., 1991). The GPI method places an evenly spaced grid over the site where a two transects are sampled at each grid point to assess disturbance from skid roads and random skid trails. The Traverse method is used for measuring infrequent types of disturbance such as skids, fire guards, and haul roads.

A study carried out by McMahon (1995), looked at the accuracy of two ground survey methods for assessing site disturbance. McMahon outlined three ground survey methods, which have been used by researchers to assess site disturbance to be Point Transect (PT) method, Line Transect (LT) method, and Grid Point Intercept (GPI) method. He also mentioned a study carried out by Firth et al. (1984) where aerial photographs combined with ground reconnaissance were used to assess site disturbance. The study found this method had several advantages over the ground-based methods, such as rapidly assessing large areas for deep disturbance. However, identifying less severe disturbance and disturbance without a distinct colour difference were difficult to identify and required more ground surveying to be identified.

3.2 Satellite imagery

Satellite imagery is a readily available resource with around 95% of New Zealand being mapped and accessible through Land Information New Zealand (Land Information New Zealand, 2020). Another platform that provides free satellite imagery is Google Earth (Google Earth, 2020). If desired, more detailed imagery can be carried out using the desired imagery option such as drones, fixed wing aircrafts or satellites to help provide more current and/or higher resolution images. These images have many uses, such as determining land cover of an area, looking at land use change over time or providing weather information.

Planet.com uses three different satellite constellations, PLANETSCOPE, RAPIDEYE, and SKYSAT. The constellation used to take the images for this study will be the PLANETSCOPE constellation, they have a pixel size of 3 by 3 meters. The multispectral sensors aboard the satellites collect information from different wavelengths like a digital camera, however, there are two main differences. Firstly, instead of limiting itself to visible wavelengths it detects a much broader range of wavelengths not visible to the human eye, such as infrared and sometimes thermal. Secondly, instead of automatically combining the information from each wavelength to create a picture, information for each specific wavelength range is stored as a separate image, this is commonly called a 'band' (Horning, 2004). The PLANETSCOPE has four bands which include Red Blue and Green (RGB), and a near infrared band (NIR). These bands when viewed alone are like a black and white photograph, it is up to the user to combine the images from different wavelengths to create the desired colour image. The combination of bands can be done to highlight certain features within an image such as vegetation, water, or soil. The following list provides information on the first four bands and was retrieved from (Horning, 2004). It provides a generalised wavelength range and common uses for each band.

Band 1 (0.45-0.52 μm , blue-green): Since this short wavelength of light penetrates better than the other bands it is often the band of choice for aquatic ecosystems. It is used to monitor sediment in water, mapping coral reefs, and water depth. Unfortunately, this is the noisiest of the Landsat bands since short wavelength blue light is scattered more than the other bands. For this reason, it is rarely used for "pretty picture" type images.

Band 2 (0.52-0.60 μm , green): This has similar qualities to band 1 but not as extreme. The band was selected because it matches the wavelength for the green we see when looking at vegetation.

Band 3 (0.63-0.69 μm , red): Since vegetation absorbs nearly all red light (it is sometimes called the chlorophyll absorption band) this band can be useful for distinguishing between vegetation and soil and in monitoring vegetation health.

Band 4 (0.76-0.90 μm , near infrared): Since water absorbs nearly all light at this wavelength, water bodies appear very dark. This contrasts with bright reflectance for soil and vegetation, so it is a good band for defining the water/land interface.

3.3 Image processing

Images can be processed using a range of tools available through different programs. Image classification is a very useful tool for determining the land cover of large areas. Three of the main types of image classification are supervised, unsupervised and object-based image analysis (OBIA). Supervised classification uses training samples to classify the image while the unsupervised classification finds spectral classes without the analyst's intervention. OBIA groups pixels into representative shapes with size and geometry. When using low spatial resolution all three techniques work well, but when you have high spatial resolution, OBIA is superior to traditional pixel based classification (GISGeography, 2014). For the low spatial resolution, supervised classification outperforms unsupervised.

Traditional remote sensing data classification methods include Maximum Likelihood Estimation classifier, distance measure, clustering, or logistic regression (Rana et al., 2020). Maximum Likelihood Estimation (MLE) is one of the most popular supervised classification methods used with remote sensing image data (Mustapha, 2010). When using this tool through ArcMap 10.7.1 it uses an algorithm that is based on two principles, 1) The cells in each sample in the multidimensional space being normally distributed and 2) Bayes theorem of decision making (ESRI, 2019a).

Cells in a class are rarely homogeneous. For example, if irrigated land in the shade has a reflectance that is similar to non-irrigated land in the sun they could end up in the same class. To classify pixels that fall within the overlap of the two classes, the MLE classifier calculates the probability of the cell belonging to each class then it is assigned to the class with the higher probability, resulting in the term "maximum likelihood" (ESRI, 2019c).

Over the last decade, more advanced methods, such as decision trees, k-nearest-neighbours, random forest, neural networks and Support Vector Machine (SVM) have been used for Land Use and Land Cover (LULC) mapping (Runa et al., 2020). A recent study carried out in 2016 by Khatami et al., compared state of the art supervised classification methods for land cover classifications. It was found that the more recent and advanced classifiers, such as SVM, k-nearest-neighbours, and random forest tree generally provide better results than the traditional classifiers, with SVM being the most efficient method (Khatami et al., 2016)

SVM and Random Forest Tree classifiers are recent developments in the computational aspects of image classification. Their ability to minimize classification errors make them superior to the parametric classifiers such as MLE classifier. The basic approach in SVM is to identify a hyper-plane that produces an optimal separation between the two classes (Runa et al., 2020). This classifier does not assume that the cells in each class are normally distributed.

A major limitation when using supervised classification techniques for mapping land cover is that they are designed to create mutually exclusive classes, with each pixel being pure and only comprising of a single class. This is most often not the case, pixels containing mixed land cover are usually abundant, particularly near imprecise and fuzzy boundaries. These mixed pixels are a major problem for

supervised classification as the output is “hard”, containing only the allocated class, therefore, these techniques cannot represent mixed pixels limiting the accuracy of the classification (Foody, 1996).

Assessing the accuracy of a classification is an essential step after performing a land cover classification to assess the errors of the map and its reliability. There are many methods for assessing the accuracy of a classification, the most widely prompted and used, however, may be derived from a confusion matrix, it is currently at the core of accuracy assessment literature (Foody, 2001).

One of the most common methods for assessing the accuracy of the classified image is to create a set of random points over the study area to collect ground-truth data from and compare the results using a confusion matrix. This allows for comparison between the classified image and another data source that is considered to be accurate or ground truth data (ESRI, 2019e). Ground-truth data can be either collected out in the field, however, this is time consuming and expensive, or it can be derived from interpreting existing classified imagery or high-resolution imagery and GIS data layers.

4. Problem statement and Objective

Forestry soil disturbance assessments are generally carried out using ground-based methods, such as plots and line transects. These methods are both labour intensive and time consuming. With the increasing availability of daily satellite imagery and improving image classification tools, there is the opportunity to use these resources to change the way site disturbance is measured.

- Determine if satellite imagery and supervised image classification can be used to accurately identify bare soil on harvested forestry sites.

5. Methodology

5.1 Site locations

A total of seven sites were analysed throughout this study. Three sites were from the Golden Downs area near Nelson and are managed by Onefortyone. The three locations can be seen below in *Figure 1*. These sites ranged in size from 71 ha to 330 ha and had all been recently harvested with some areas that had undergone mechanical site preparation.



Figure 1: Location of harvest sites near Nelson.

Four sites with proximity to Christchurch city were used, which are managed by Laurie Forestry. The locations of these four sites near Christchurch can be seen below in *Figure 2*. These four sites range in size from 2 ha to 10 ha and have all been mechanically prepped and recently planted.

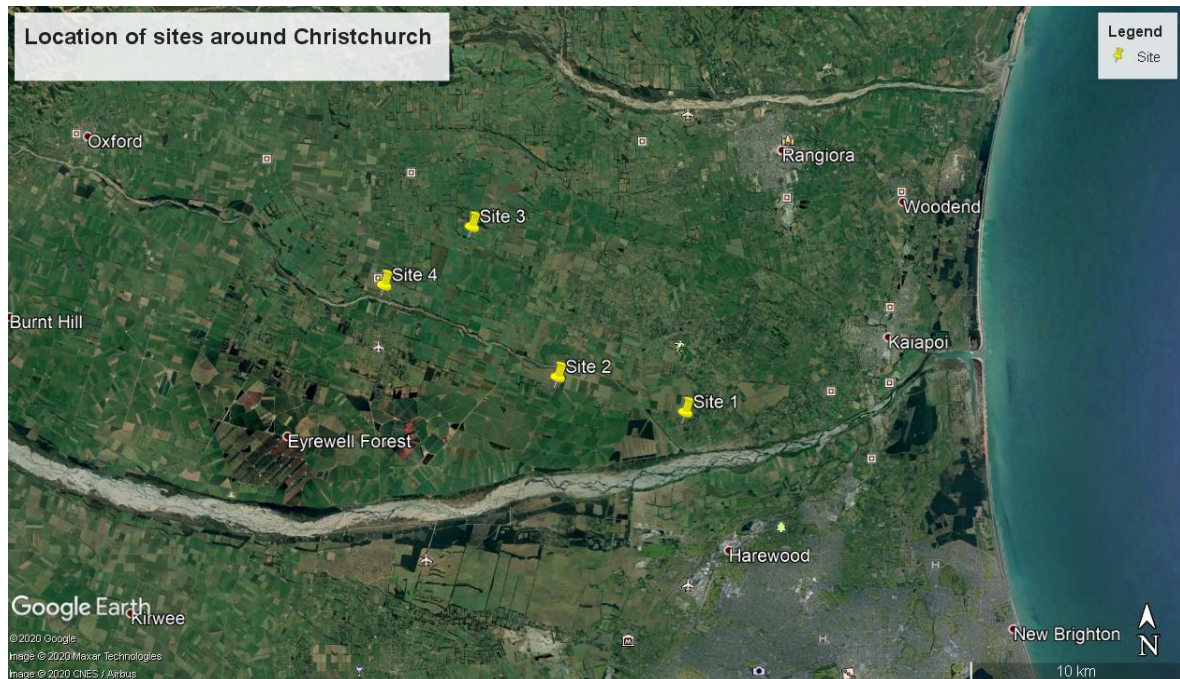


Figure 2: Location of four harvest sites around Christchurch.

Satellite images of each harvest site were taken by the PLANETSCOPE constellation, the images of the Nelson harvest areas were provided by OneFortyOne and the areas around Christchurch were accessed through a Planet.com Education & Research basic account, which allowed access to daily imagery of the harvest sites. Each harvest area was given a corresponding number to keep track of each site during analysis. The satellite images for each site were uploaded into ArcMap 10.7.1 Desktop software and projected using the WGS84 coordinate system, this corrects any errors associated with different projections to help improve accuracy between sites, and is the same coordinated system used when carrying out ground truthing.

Figure 3 shows an original satellite image of a harvest area taken from the PLANETSCOPE constellation. Using the polygon tool on ArcMap, the harvest boundary is outlined in red as seen in *Figure 4*. This allows for the pixels within the harvest boundary to be processed instead of processing the entire image, this is done using the 'clip' tool on ArcMap which uses the previously created polygon (harvest boundary) to extract all the pixels within the polygon and create a new image of the harvested area of interest.



Figure 3: Original satellite image of harvest area.



Figure 4: Outline of harvest area.

5.2 Ground-truth data

An onsite analysis was carried out for each harvest site to create ground-truth plots. Due to time constraints and the large geographical area of some of the sites, random plots covering the entire site were not a feasible option, instead at each site two separate lines of 15 points were created with 10 metre spacings between points. At each point, the location was recorded using a Trimble Geo7X Handheld together with a Trimble Zephyr 3 GPS receiver, the system was held steady and level at each location for at least 60 seconds which provided an accuracy of 10-15 centimetres.

At each point, a visual inspection was carried out to determine if the major landcover type was bare soil or not bare soil (e.g. slash or vegetation) for a 3 metre radius around the point, this was recorded along with each point number. If the landcover was determined to be slash, this was then further broken down into 5 sub classes using a 1 – 5 scale, 1 being light slash cover and 5 being heavy slash cover. To supplement the visual inspection and for later analyses, a photo was taken at each point of the land cover with the point number and observed land cover type. *Figure 5* provides an example for some of the possible land cover types, it shows a light and heavy slash example (1 and 5 on the scale range respectively), as well as a bare soil and a vegetation example. The light slash, heavy slash, and vegetation are all classified as ‘Slash’ while the bare soil is the only one classified as ‘Bare soil’.



Figure 5: Example of some landcover types.

Once all the ground truth data was collected it was later transferred into an Excel file where each point was given a value of either 'S' for slash or 'B' for bare soil, if it was slash it also had an associated column with a number for the amount of slash cover using the 1-5 scale mentioned above.

The ground-truth plots were differentially corrected using the known location of a base station, and individually correcting each point at the same time that point was created, the points were then uploaded into ArcMap and projected using the same WSG84 coordinate system as the original satellite images. *Figure 6* provides an example of a harvest site with the ground-truth points overlaid, the ground-truth plots are the yellow dots that can be seen on the lower and upper, right-hand side of the harvest area.



Figure 6: Ground-truth plots over satellite image.

5.3 Analysing satellite images

The images taken by the PLANESCOPE constellation are made up of four spectral bands, which represent different wavelengths and can be arranged to suit the individual study as outlined in the literature review. For this study, the near infrared band is used along with the 2nd and 3rd band. This step is carried out by opening up the properties of the 'clipped' image going to the symbology tab and assigning the Red Green and Blue colours to the desired bands, in this case the near infrared (NIR) or 4th band was assigned to Red, the 2nd band was assigned to Green, and the 3rd band was assigned to Blue. *Figure 7* shows the rearranged bands of the harvest site which produce a near infrared image. The light to dark red areas seen in *Figure 7* are slash and standing vegetation, while the light blue and white areas are bare soil.

From here, the colours of the image can then be inverted to provide a more distinguishable contrast between the two classes, this is also done under the symbology tab by checking the 'invert' box. *Figure 8* shows how this changes the colours compared to the infrared image in *Figure 7*. Shown in *Figure 8*, the dark brown areas are bare soil and the white to blue areas are slash & vegetation. It should be noted that the shaded areas of the harvest site as seen in the left corner of the harvest area appear dark in *Figure 7* and bright in *Figure 8*, these dark areas are the shaded parts of the harvest site. Distinguishing between the two classes is difficult in these shaded areas as only large areas of bare soil are visible, otherwise the areas have a reflectance similar to slash.



Figure 7: Rearranged bands of harvest area.



Figure 8: Inverted bands of harvest area.

5.4 Creating and evaluating training samples

Selecting appropriate training samples is the most important step in the image classification process as they are used to train the software. Using the image in *Figure 8*, training samples can be created for the two classes either bare soil or slash & vegetation. This is done by zooming into sections of the image where each class can be easily identified. An example of a training sample for bare soil and for slash can be seen in *Figure 9 and 10*, respectively. The sample in *Figure 9* is part of a skid site and contains pixels of pure bare soil, a road is visible on the right-hand side of *Figure 10* which could also be used as a bare soil training sample. The sample in *Figure 10* includes areas of both vegetation and slash, the darker blue areas are vegetation and the light blue and white areas are slash.




Figure 9: Training sample of bare soil.




Figure 10: Training sample of slash and vegetation.

Training samples can be created using a range of different shapes or custom polygons, in these examples I have used a square to identify areas of bare soil or not bare soil, the training sample collection process is a matter of manual interpretation between the two classes where the location and size of the training sample comes down to the analyst's interpretation. When selecting training samples, they should be evenly spread over the site, aiming to cover the entire spectral range and variability for each class. Training samples should not include pixels where the ground-truthing was carried out as these points are later used to assess the accuracy of the classification tool.

Once the required number of training samples have been created, they can be evaluated to ensure they will provide an accurate classification. This is done by displaying Histograms, Scatter plots, and the Statistics for the band ranges within the group of training samples. Each should be examined so that there is minimal to no overlap between the two classes. If there is overlap, this means some training samples need to be removed or redone to minimise the overlap. This is an iterative process and should be repeated until the user is happy with the training sample set. An example of the Histograms, Scatter plots, and Statistics can be seen in *Figure 11, 12 and 13*, respectively.




Figure 11: Histograms showing band range for soil and slash.

Figure 11 shows Histograms with the spectral range for each class and the associated count of individual pixels over that range, where red represents the bare soil class and green represents the not bare soil class. *Figure 12* shows the relationship between the individual bands and makes it easy to identify areas of overlapping between the two classes. *Figure 13* provides a table of statistics for all

bands within both classes, this also makes it easy to find overlapping classes by looking at the minimum and maximum values for each band between the two classes. Some of the training samples used in this example had to be redone to minimise the overlap present in the three figures.




Figure 12: Scatter plots showing band range for soil and slash.

Statistics				
Bare Soil				
Statistics	Band_1	Band_2	Band_3	Band_4
Minimum	5031.00	3996.00	2999.00	4105.00
Maximum	9952.00	10459.00	10995.00	9190.00
Mean	8041.83	8167.33	7818.31	6732.28
Std.dev	982.21	1407.30	1756.72	1201.98
Covariance				
Band_1	964737.89	1206100.53	1350752.50	752106.98
Band_2	1206100.53	1980483.77	2407263.99	1431418.13
Band_3	1350752.50	2407263.99	3086063.83	1903088.87
Band_4	752106.98	1431418.13	1903088.87	1444751.71
Slash/V...				
Statistics	Band_1	Band_2	Band_3	Band_4
Minimum	4283.00	3106.00	1832.00	3435.00
Maximum	6451.00	5401.00	5341.00	7640.00
Mean	5154.93	4172.50	3554.40	5450.37
Std.dev	363.88	393.15	709.58	859.50
Covariance				
Band_1	132406.09	133905.31	221261.50	-99225.50
Band_2	133905.31	154570.74	255253.35	-19737.96
Band_3	221261.50	255253.35	503504.44	-40841.82
Band_4	-99225.50	-19737.96	-40841.82	738739.20

Figure 13: Statistics for band ranges of soil and slash.

Once a good set of training samples is attained, they can be saved as a signature file. This signature file can then be used to create both a Support Vector Machine classifier file and a Maximum Likelihood Estimation classifier file, these files can be used to create the final classification for that set of training samples.

5.5 Training sample sensitivity analysis

A sensitivity analysis was carried out on training sample size, this was done by repeating the above steps for training samples for each class, the number of samples was increased in intervals of 5 for each class, ranging from 5 training samples per class to 20 training samples per class. Each varying set of training samples were evaluated using the Histograms, Scatter plots, and Statistics as seen above to ensure they were as accurate as possible. The following figure provides an example of the varying training sample sets and how they were spread out over the site to capture the entire spectral range and variability of each class. The left image in *Figure 14* represents 5 training sample per class, while the right image represents 20 training sample per class.




Figure 14: Comparison of training sample sets.

Figure 15 and 16 show the results for a classified image using the 5 and 20 training samples per class, that were shown above.



Figure 15: Classified image using 5 training samples per class.




Figure 16: Classified image using 20 training samples per class.

Once all the images have been classified, the ground truth points can be used to assess the accuracy of the classification. Figure 17 provides a close up of the ground-truth plots over a classified image, from here each pixel that has a ground-truth value associated with it can be assessed to determine if it has been correctly classified according to the ground-truth results. This is done by adding a column

for each training sample size to the Excel file which states the classified value of each pixel using the same code as the one used for the ground truth values, 'S' for slash and 'B' for bare soil.




Figure 17: Close up of ground-truth points

For each site, an Excel spread sheet was created that contained the following columns seen in *Figure 18*. The first column contains the point number, the second is the ground truth value associated with that point either 'S' or 'B', and the slash cover value associated with that plot (1 – 5) if it was classed as slash. The rest contain the classified values for each pixel associated with each plot again being either 'S' or 'B'. This example shows the values for the SVM classification, hence, the SVM and the number represents the number of training samples used for each class. Another table was created similar to this for each site except for the MLE.

Plot	Ground-truth	Slash Cover (1-5)	SVM 5 samp	SVM 10 samp	SVM 15 samp	SVM 20 samp
------	--------------	-------------------	------------	-------------	-------------	-------------

Figure 18: Spread sheet example

After all the columns are filled out with the ground-truth and classified data, a confusion matrix can be computed for each training sample set. This is done by comparing the ground-truth values to those in the varying training sample columns to determine how many agreed and how many disagreed. The main inputs into the confusion matrix are the numbers highlighted in grey in *Table 1*. The number 5 for slash under the ground-truth column represents the number of slash pixels that were correctly classified as slash, while the 1 represents slash pixels that were incorrectly classified as bare soil. The 6 under bare soil in the ground-truth column represents bare soil pixels that were incorrectly classified as slash, while the 18 represents bare soil pixels that were correctly classified as bare soil.

Table 1: Example of confusion matrix and Cohen's Kappa.

SVM				
Cohen's Kappa				
Site 2: 5 Samples				
	Ground-truth			
Classified	Slash	Bare Soil	Total	User's
Slash	5	6	11	0.45
Bare Soil	1	18	19	0.05
Total	6	24	30	
Producer's	0.83	0.25		0.77
Agreement	5	18	23	
By Chance	2.2	15.2	17.4	
Kappa	0.44			

A confusion matrix was computed for each site, training sample size, and classification method to help analyses the results. Each matrix includes errors of omission and commission, overall accuracy and derives a Kappa index of agreement between the classified image and the ground truth data.

The errors of commission, known as the User's accuracy or type 1 error, are false positives where pixels are incorrectly classified as a known class when they should be classified as something else. For example, the classified image identifies the pixel as bare soil when the ground truth identifies it as slash. For this case, the bare soil class has extra pixels than it should according to the ground truth data.

The errors of omission known as the Producer's accuracy, or type 2 error, are false negatives where pixels of a known class are classified as something else. For example, the classified image identifies a pixel as slash, but it should be bare soil. And for this case the bare soil class is missing pixels according to the ground truth data.

The overall accuracy of the classification is the total number of pixels that agree for all classes in the classification, using the matrix in *Table 1*, for example, 5 for slash and 18 for bare soil, these are the diagonal values of the matrix and can also be seen in the agreement row with a total of 23 pixels agreeing. This value is divided by the total number of pixels that are being examined (30), producing a total agreement of 77 % for this classification as seen in *Table 1*.

The table also provides a Kappa index of agreement, which provides an overall assessment of the accuracy of the classification but takes into consideration the agreement due to chance and subtracts it from the overall agreement. In *Table 1* Slash represents $6/30 = 20\%$ of the ground-truth values and $11/30 = 37\%$ of the classified values. Thus $20\% * 37\% = 7.3\%$ of the agreement between the classified and ground-truth values for Slash are by chance i.e. $7.3\% * 30 = 2.2$ of the pixels. Similarly, 15.2 of the bare soil pixels are by chance, which means the total number of pixels classified by chance is 17.4. Now subtracting out the agreement due to chance using equation (1) will give the final Kappa index of agreement of 0.44 (44%) as seen in *Table 1*. Kappa values range from 0 – 1 where zero is no agreement and 1 is total agreement, sometimes a negative value is returned which also indicates no agreement.

$$k = \frac{n_a - n_e}{n - n_e} \quad (1)$$

Where:

k = Kappa Index

n = Number of pixels

n_a = Number of agreements

n_e = Number of agreements due to chance

There is no universal agreement for a set of Kappa values and how significant the agreement is, therefore the following table provided by McHugh (2017) will be used in the study as it is similar to a range of other values provided in literature.

Table 2: Ranges for Kappa agreement (McHugh, 2017).

Value of Kappa	Level of Agreement	% of Data that are Reliable
0-.20	None	0-4%
.21-.39	Minimal	4-15%
.40-.59	Weak	15-35%
.60-.79	Moderate	35-63%
.80-.90	Strong	64-81%
Above .90	Almost Perfect	82-100%

6. Results and Data analysis

Table 3 and 4 show the results for the overall agreement and the Kappa index for the varying sample size. Table 3 shows the results when using the MLE classifier, while Table 4 shows the results when using the SVM classifier. The numbers highlighted in yellow show the best results for each site, for each classifier, and at which training sample size it occurred.

Table 3: Overall accuracy and Kappa index for Maximum likelihood classification.

		5 Samples		10 Samples		15 Samples		20 Samples	
		Overall	Kappa	Overall	Kappa	Overall	Kappa	Overall	Kappa
MLE	Christchurch								
	Site 1	0.7	0.31	0.6	0.15	0.67	0.25	0.63	0.2
	Site 2	0.67	0.31	0.57	0.25	0.57	0.2	0.63	0.27
	Site 3	0.5	-0.06	0.53	0.05	0.5	0	0.53	0.05
	Site 4	0.33	-0.01	0.57	-0.05	0.57	-0.15	0.57	-0.15
Nelson	Site 1	0.53	0.09	0.53	0.09	0.5	0.03	0.53	0.1
	Site 2	0.83	0.38	0.87	0.27	0.87	0.27	0.9	0.35
	Site 3	0.83	0.44	0.83	0.51	0.77	0.44	0.77	0.44

Table 4: Overall accuracy and Kappa index for Support Machine Vector classification.

		5 Samples		10 Samples		15 Samples		20 Samples	
SVM		Overall	Kappa	Overall	Kappa	Overall	Kappa	Overall	Kappa
Christchurch	Site 1	0.77	0.43	0.5	-0.03	0.67	0.25	0.57	0.11
	Site 2	0.77	0.44	0.73	0.39	0.77	0.39	0.77	0.39
	Site 3	0.57	0.18	0.6	0.23	0.57	0.2	0.53	0.15
	Site 4	0.53	-0.3	0.57	-0.15	0.57	-0.15	0.57	-0.27
Nelson	Site 1	0.57	0.18	0.53	0.06	0.37	-0.21	0.3	-0.34
	Site 2	0.57	0.02	0.87	0.27	0.97	0.65	0.97	0.65
	Site 3	0.73	-0.11	0.73	-0.11	0.77	0.1	0.63	-0.08

From these two tables, it is clear that the SVM classifier provides a more accurate classification than the MLE classifier in terms of overall accuracy and Kappa index when looking at the best-case scenarios, except for site 3 in Nelson. It is also clear that the Christchurch sites are most accurate with a smaller training sample sets compared to the larger Nelson sites.

Taking the average results for both overall agreement and Kappa index for all training sample sizes returns an overall agreement of 64% and Kappa index of 0.19 for the MLE classifier and an overall agreement of 65% and Kappa index of 0.18 for the SVM classifier. When taking the average Kappa index negative values were treated as zero as they mean no agreement and including the negative values throws off the average Kappa values. If we take the best-case scenarios from each site the MLE classifier returns an average overall agreement of 68% and an average Kappa index of 0.23, while the SVM classifier returns an average overall agreement of 72% and an average Kappa coefficient of 0.27.

The following two tables provide the Producer's and User's accuracy for the varying training sample size for each site. *Table 5* shows the User's and Producer's accuracy when using the MLE classifier, while *Table 6* shows the User's and Producer's accuracy when using the SVM classifier.

Table 5: Producer's and User's accuracy when using Maximum Likelihood classification.

			5 Samples		10 Samples		15 Samples		20 Samples	
MLE			Producer's	User's	Producer's	User's	Producer's	User's	Producer's	User's
Christchurch	Site 1	Slash	0.56	0.5	0.56	0.38	0.56	0.45	0.56	0.42
		Disturbed	0.24	0.2	0.38	0.24	0.29	0.21	0.33	0.22
	Site 2	Slash	0.83	0.36	1	0.32	0.83	0.29	0.83	0.33
		Disturbed	0.38	0.06	0.54	0	0.5	0.08	0.42	0.07
	Site 3	Slash	0.61	0.58	0.56	0.63	0.5	0.6	0.56	0.63
		Disturbed	0.67	0.64	0.5	0.57	0.5	0.6	0.5	0.57
	Site 4	Slash	0.57	0.19	0.29	0.2	0.14	0.13	0.14	0.13
		Disturbed	0.74	0.33	0.35	0.25	0.3	0.27	0.3	0.27
Nelson	Site 1	Slash	0.79	0.5	0.79	0.5	0.79	0.48	0.86	0.5
		Disturbed	0.69	0.38	0.69	0.38	0.75	0.43	0.75	0.33
	Site 2	Slash	0.82	1	0.89	0.96	0.89	0.96	0.93	0.96
		Disturbed	0	0.71	0.5	0.75	0.5	0.75	0.5	0.67
	Site 3	Slash	0.92	0.88	0.88	0.91	0.75	0.95	0.75	0.95
		Disturbed	0.5	0.4	0.33	0.43	0.17	0.55	0.17	0.55

Table 6: Producer's and User's accuracy when using Support Vector Machine classification.

			5 Samples		10 Samples		15 Samples		20 Samples	
SVM			Producer's	User's	Producer's	User's	Producer's	User's	Producer's	User's
Christchurch	Site 1	Slash	0.56	0.63	0.44	0.29	0.56	0.45	0.56	0.38
		Disturbed	0.14	0.18	0.48	0.31	0.29	0.21	0.43	0.25
	Site 2	Slash	0.83	0.45	0.83	0.42	0.67	0.44	0.67	0.44
		Disturbed	0.25	0.05	0.29	0.06	0.21	0.09	0.21	0.09
	Site 3	Slash	0.44	0.73	0.5	0.75	0.39	0.78	0.33	0.75
		Disturbed	0.25	0.53	0.25	0.5	0.17	0.52	0.17	0.55
	Site 4	Slash	0	0	0.14	0.13	0.14	0.13	0	0
		Disturbed	0.3	0.3	0.3	0.27	0.3	0.27	0.26	0.29
Nelson	Site 1	Slash	1	0.52	0.5	0.5	0.71	0.4	0.64	0.36
		Disturbed	0.81	0	0.44	0.44	0.94	0.8	1	1
	Site 2	Slash	0.57	0.94	0.89	0.96	1	0.97	1	0.97
		Disturbed	0.5	0.92	0.5	0.75	0.5	0	0.5	0
	Site 3	Slash	0.92	0.79	0.92	0.79	0.92	0.81	0.75	0.78
		Disturbed	1	1	1	1	0.83	0.67	0.83	0.86

In general, when comparing *Table 5 and 6* above, the producer's accuracy is generally higher when using MLE classifier, while the user's accuracy is higher when using the SVM classifier.

7. Discussion

The study carried out in 2016 by Khatami et al. compared more recent supervised image classification methods to the more traditional supervised classification methods such as Maximum Likelihood. It was found that the more recent methods generally outperform the traditional methods with the Support Machine Vector classifier generally providing the best results. From *Table 3 and 4* I have confirmed that Support Machine Vector classifier provides more accurate results compared to the Maximum Likelihood classifier when used on harvested forestry sites.

The SVM classifier does not assume the cells within a class are normally distributed but the MLE classifier does. Looking at the histograms in *Figure 11* the cells in the different classes are not always normally distributed. This is due to the 'Slash' class not only including slash but also vegetation, and the possibility of mixed pixels in both classes, this makes the SVM classifier more suited for classifying harvest sites.

When assessing the accuracy of the classified images the method that is most widely prompted and used, may be derived from a confusion matrix, it is currently at the core of accuracy assessment literature (Foody, 2001). When collecting the ground-truth data to help populate the confusion matrix, the most common method is creating a set of random points over the study area and collecting ground truth data at these points (ESRI, 2019e). The method used in this study consisted of a total of 30 ground-truth points, split in to two lines of 15 points. This method provided a much faster way to collect the ground-truth data, rather than covering the entire harvest area going to randomly scattered points. However, this method also has its limitations as it doesn't include as many assessment points as a random scatter which would provide a truer assessment of the accuracy, it could also miss variability within the sites.

When using the best classifications from each site for each classification method, the MLE and SVM return relatively high overall agreement results of 68% and 72% respectively, with the ground-truth data. However, when looking at the Kappa coefficient for both the classification methods they returned average results of 0.23 and 0.27, according to *Table 2* both classifiers fall in the category of minimal agreement where only 4% - 15% of the data is considered to be reliable.

For the larger Nelson sites as the number of training samples increase so did the accuracy of the classification, while the smaller sites around Christchurch have the best results when only using 5 or 10 training samples. With the larger sites there is a higher chance that there will be more variability throughout the site which will require more training samples to accommodate for this variability when performing the classification. While the smaller sites required less samples to cover the variability throughout the site and including more training samples increases the chance of mixed pixels being included when training the classifier.

The most accurate classification was Site 2 in Nelson and returned an overall accuracy of 97% and a Kappa coefficient of 0.65 which is only moderate agreement according to the Kappa values in *Table 2* but an almost perfect overall accuracy. This was achieved using the SVM classifier. The worst classification was Site 1 in Nelson with an overall accuracy of 30% and a Kappa index of -0.34 meaning there was no agreement according to Kappa index and poor overall accuracy, this was also attained using the SVM classifier.

When comparing the Producer's and User's accuracy for each site the general trend is that the Producer's accuracy is higher than the User's. This means that map accuracy from the map producer's point of view is greater than the accuracy from the map user's point of view. For example looking at the best classification, Site 2 from Nelson using 15 training samples has a 100% producer's accuracy for slash meaning 100% of the reference slash areas have been correctly identified as slash, but has a user's accuracy of 97% meaning only 97% of the areas identified as slash in the classified image were actually slash.

Also, if you look at the Producer's and User's accuracy between the two different classification methods, the Producer's accuracy is higher when using the MLE classifier, however the User's accuracy is higher when using the SVM classifier. This means that the maps created using the SVM classifier are more accurate from the map users' point of view.

Looking at the following two figures is it easy to identify the large deep disturbances such as roads, skid sites and large skid trails in the original satellite image in *Figure 19*. These features are easily picked up in the classification as they have the greatest contrast compared to the slash class and can be easily identified from the original images as seen in *Figure 20*. The areas with lighter disturbances proved difficult to identify.



Figure 19: An original satellite image of a harvest site.




Figure 20: Classified satellite image using the SVM classifier.

Pixels that are being incorrectly classified are generally close to a transition zone between the two classes and/or they contain a bit of each class, these pixels are known as mixed pixels. For example, a pixel that contains slash but has been mixed with soil will be difficult to classify and may result in misclassification. These mixed pixels are a major problem for supervised classification as the output is “hard”, containing only the allocated class, therefore these techniques cannot represent mixed pixels limiting the accuracy of the classification (Foody, 1996).

Another reason for misclassification of pixels and a major limitation for steeper sites is the effect caused by shaded areas as it changes the reflectance of bare soil making it appear as slash, this can be seen when comparing *Figure 19* and *20* when looking at the shaded areas at the top of the site in *Figure 19* and how these areas are being classified as slash in *Figure 20*, e.g. the large green areas. This effect can be minimised by finding imagery that was taken mid-day or when there is the least amount of shading on site as possible. The shading can have an impact not just on hill slopes but also slash piles, wind rows, and from other larger standing vegetation within or around the site.

When an area was identified as slash, it had an associated slash cover value ranging from 1-5 where 1 was light slash and 5 was heavy slash. The sale was put in place to determine if the level of slash had a direct impact on the pixel being correctly classified as slash, ground-truth values that were recorded as heavy slash were expected to almost certainly be classified as ‘Slash’ compared to light slash values. This was not the case and no trends were found between the slash value and the classified pixel.

As higher resolution images become more readily available it will make classifications more accurate by decreasing the pixel size, it will increase the number of ‘pure’ pixels (pixels that contain only one class). To get a better gauge on the accuracy of the classifications ideally, a sample of random ground-truth plots over each site would need to be carried out.

8. Conclusion

This study aimed to determine if satellite images could be used in conjunction with supervised image classification to accurately identify bare soil on harvested sites. A total of seven harvested sites were analysed, three of which were in Nelson, while the other four were around Christchurch. Ground-truth data was collected at each site where a total of 30 points were analysed, the 30 points were split into two lines of 15 points spaced at 10 metre intervals over the harvest site. At each point the geographical location was recorded along with the dominant land cover value, either bare soil or not bare soil (e.g. slash or vegetation).

Satellite images were used that were taken as close to the date the ground-truthing was carried out while also taking into consideration the time of day the image was taken and the quality of the image. These images could then be classified using two supervised classification methods, either the MLE classifier or the SVM classifier. The classified images were compared to the ground-truth points using a confusion matrix.

The average overall agreement for MLE and SVM was 64% and 65% respectively, while the Kappa index was 0.19 and 0.18, respectively. However, when looking at just the best-case scenarios for each classification technique the average overall agreement for MLE and SVM was 68% and 72% respectively, while the Kappa index was 0.23 and 0.27, respectively. Confirming that the SVM classifier generally outperforms the MLE classifier. Although the overall accuracy was high for both the average and best-case scenarios, the Kappa index for both cases returned values of low agreement.

When looking at Producer's and User's accuracy for the two classifiers the SVM generally returned a higher user's accuracy, while the MLE generally returned a higher producer's accuracy meaning the SVM classified map would be more accurate from the users' point of view.

Both classification techniques easily pick up large disturbances such as skids sites and roads but struggle with smaller lighter disturbances where pixels contain some of both classes, for example, the pixel may be mainly disturbed but also contain some heavy slash, this is called a 'mixed' pixel and can be difficult to classify, reducing the accuracy of the classification. Another factor affecting the accuracy of the classifications is shaded areas, the shade changes the reflectance of the ground resulting in the misclassification of pixels.

As higher resolution images become more readily available it will make classifications more accurate by decreasing the pixel size, it will increase the number of 'pure' pixels (pixels that contain only one class) and reduce the area of mixed pixels at the boundaries between classes.

Overall, the two classification methods both have relatively high overall agreement with the ground-truth data, however, they both have low Kappa coefficients values that fall into the range of minimal agreement between the classification and ground-truth data. To get a better gauge on the accuracy of the classifications ideally, a sample of random ground-truth plots over each site would need to be carried out. This would help get ground-truth data that would cover the entire geographical area of the site as well as the variability within the site.

9. References

Curran, M. (1991). *Measuring soil disturbance following timber harvesting*. BC Ministry of Forests.

Elliot, W.J., Page-Dumroese, D., & Robichaud, P.R. (1999) *The Effects of Forest Management on Erosion and Soil Productivity*. In L. Rattan (Ed.), *Soil quality and soil erosion* (195- 208). Iowa. Soil and Water Conservation Society.

a) ESRI. (2019). *Evaluating training samples*:

<https://desktop.arcgis.com/en/arcmap/latest/extensions/spatial-analyst/image-classification/evaluating-training-samples.htm>.

b) ESRI. (2019). *How Maximum Likelihood Classification works*:

<https://desktop.arcgis.com/en/arcmap/10.3/tools/spatial-analyst-toolbox/how-maximum-likelihood-classification-works.htm>.

c) ESRI. (2019). *Performing the classification*:

<https://desktop.arcgis.com/en/arcmap/10.7/tools/spatial-analyst-toolbox/performing-the-classification.htm>.

d) ESRI. (2019). *Train Support Vector Machine Classifier*:

<https://desktop.arcgis.com/en/arcmap/10.3/tools/spatial-analyst-toolbox/train-support-vector-machine-classifier.htm>

e) ESRI. (2019). *Accuracy assessment for Image Classification*:

<https://desktop.arcgis.com/en/arcmap/latest/manage-data/raster-and-images/accuracy-assessment-for-image-classification.htm>

Firth, J., Van Dijk, W. A. J., & Murphy, G. (1985). *A preliminary study of techniques for estimating harvesting-related soil disturbances from aerial photographs*. Forest Research Institute, New Zealand Forest Service.

Foody, G. M., & Arora, M. K. (1996). Incorporating mixed pixels in the training, allocation and testing stages of supervised classifications. *Pattern Recognition Letters*, 17(13), 1389-1398.

Foody, G. M. (2002). Status of land cover classification accuracy assessment. *Remote sensing of environment*, 80(1), 185-201.

GISGeography. (2014). *Image classification techniques in remote sensing*.

<https://gisgeography.com/image-classification-techniques-remote-sensing/>.

Google. (2020). *Google Earth*: <https://www.google.com/earth/>

Horning, N. (2004). Selecting the appropriate band combination for an RGB image using Landsat imagery. *Center for Biodiversity and Conservation*.

Khatami, R., Mountrakis, G., & Stehman, S. V. (2016). A meta-analysis of remote sensing research on supervised pixel-based land-cover image classification processes: General guidelines for practitioners and future research. *Remote Sensing of Environment*, 177, 89-100.

Kvamme, C. D. (2010). A comparison of ground cover and frequency estimation methods for post-harvest soil monitoring.

Land Information New Zealand (2020). *Aerial imagery*. <https://www.linz.govt.nz/data/linz-data/aerial-imagery>.

McHugh, M. L. (2012). Interrater reliability: the kappa statistic. *Biochemia medica: Biochemia medica*, 22(3), 276-282.

McMahon, S. (1995). Accuracy of two ground survey methods for assessing site disturbance. *Journal of Forest Engineering*, 6(2), 27-33.

Ministry for Primary Industries. (2020). *Forest and wood processing data*: <https://www.mpi.govt.nz/news-and-resources/open-data-and-forecasting/forestry/>.

Mustapha, M. R., Lim, H. S., & Iafri, M. M. (2010). Comparison of Neural Network and Maximum Likelihood Approaches in Image Classification. *Journal of Applied Sciences*, 10(22), 2847-2854.

Norris, J. E., Stokes, A., Mickovski, S. B., Cammeraat, E., Van Beek, R., Nicoll, B. C., & Achim, A. (Eds.). (2008). *Slope stability and erosion control: ecotechnological solutions*. Springer Science & Business Media.

Petherick, K. (2014). *Assessing the Level of Unproductive Area in Production Forestry Sites in the South Island of New Zealand*.

Phillips, C., Marden, M., & Basher, L. (2015). Forests and erosion protection: getting to the root of the matter. *NZ Journal of Forestry*, Vol. 60, No. 2.

Planet. (2020). *PLANET IMAGERY PRODUCT SPECIFICATIONS*. https://assets.planet.com/docs/Planet_Combined_Imagery_Product_Specs_letter_screen.pdf

Rana, V. K., & Suryanarayana, T. M. V. (2020). Performance evaluation of MLE, RF and SVM classification algorithms for watershed scale land use/land cover mapping using sentinel 2

bands. *Remote Sensing Applications: Society and Environment*, 19, 100351.

Resource Management (National Environmental Standards for Plantation Forestry) Regulations 2017. LI 2017/174. (2018).

<http://www.legislation.govt.nz/regulation/public/2017/0174/latest/whole.html>

Satchell, D. (2018). *Trees for steep slopes*: <http://www.nzffa.org.nz/farm-forestry-model/why-farm-forestry/trees-for-erosion-controlsoil-conservation/report-trees-for-steep-slopes/>

Visser, R., Spinelli, R., & Brown, K. (2018). Best practices for reducing harvest residues and mitigating mobilisation of harvest residues in steepland plantation forests. *Christchurch, New Zealand: School of Forestry, University of Canterbury*, 53.