Automating woody vegetation change detection at regional scales: the problem of clouds and cloud shadows

A.G. Fisher [a], *, T. Danaher [b]

[a] Joint Remote Sensing Research Program, University of Queensland, Brisbane, QLD 4072, Australia, and School of Biological, Earth and Environmental Sciences, University of New South Wales, Sydney, NSW 2052, Australia - adrian.fisher@unsw.edu.au
[b] Department of Environment, Climate Change and Water, Alstonville, NSW 2477, Australia - tim.danaher@environment.nsw.gov.au

Abstract – This paper presents a semi-automated cloud and cloud shadow masking method, developed for SPOT5 satellite data over New South Wales (NSW), Australia. As clouds are very similar to several surface features in these data, attempts at using traditional image classification techniques have not been effective. The new method uses morphological feature extraction, where marker pixels are identified using image-specific, automatically defined criteria, and mask segments are grown using the watershed from markers transformation. Manual input is only required to add and delete marker points in order to improve the masks. The method has achieved an overall accuracy of around 80 %, with most errors due to the commission rather than omission of pixels. This has proved satisfactory, and the method has been incorporated into the processing stream for woody vegetation change detection by the NSW government.

Key words: Clouds, shadows, morphological image processing

1. INTRODUCTION

Clouds are often a problem when conducting remote sensing of the earth's surface as they can locally obscure surface features and increase reflectance. Shadows cast by clouds also obscure surface features through decreasing surface reflectance. The masking of clouds and their shadows, and the subsequent exclusion of these contaminated pixels from analysis, is an important pre-processing step in many applications.

The Statewide Landcover and Trees Study (SLATS) conducted by the state governments of Queensland and New South Wales (NSW) in Australia, is an ongoing project that uses Landsat Thematic Mapper (TM), Landsat Enhanced Thematic Mapper Plus (ETM+) and SPOT5 data, to map and monitor change in woody vegetation extent annually (Danaher et al., 2010). The presence of clouds and their shadows affects the calculation of foliage projective cover (FPC), which is used to map woody vegetation, leading to false anomalies when changes in FPC are examined.

Due to the large areas being analysed, and the ongoing data acquisition, automation is important. For example, NSW covers an area greater than 800,000 km^2, and requires at least 323 SPOT5 images for a complete coverage (Figure 1). Many of the SLATS processing steps are fully automated, such as correcting for atmospheric effects, variations in the Bidirectional Reflectance Distribution Function (BRDF), and the calculation of FPC (Meier et al., 2011). The development of an accurate fully-automated cloud and cloud-shadow masking method would remove a significant delay in the processing stream, although this has not yet been achieved.
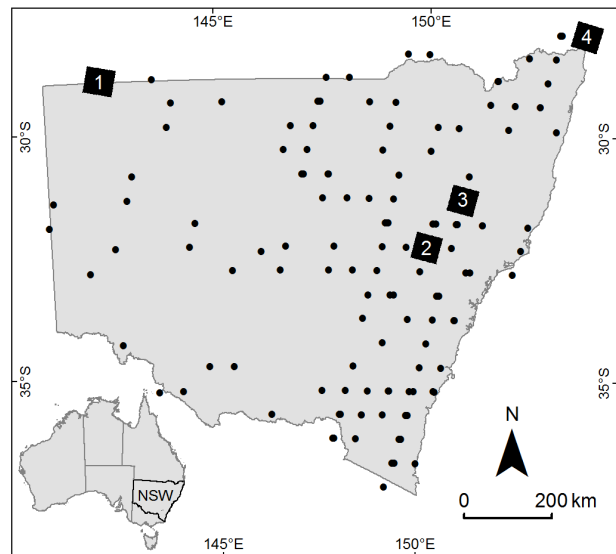

Figure 1. The location of the 142 cloudy SPOT5 images acquired during 2007-2009 over New South Wales, Australia (shown by their centre points). Subsets shown in subsequent figures were taken from the four numbered images.

2. CLOUDS IN SPOT5 DATA

The SLATS method was originally developed to use Landsat TM/ETM+ data with 30 m pixels, exploiting over 20 years of data. Detecting clouds in these data is aided by thermal infrared bands, as clouds often have a lower temperature than the ground surface (Huang et al., 2010). Improving the spatial resolution of SLATS by using SPOT5 data (10 m multispectral and 2.5 m panchromatic) has required a method that can work without a thermal infrared band. Such a method must distinguish clouds using the four available bands: green (500-590 nm), red (610-680 nm), near-infrared (780-890 nm) and mid-infrared (1580-1750 nm).

Clouds have a large variety of shapes, textures, and reflectance characteristics, and in SPOT5 data they are similar to reflective non-cloud features, such as sandy beaches or bare soil (Figure 2). Attempts at using traditional image classification techniques to detect clouds have not been effective, and usually result in the commission of these reflective surface features and the omission of thin, less reflective cloud. To avoid these errors Le Hégarat-Mascle and André (2009) developed a method that detected potential cloud objects, and removed those that did not have a matching cloud-shadow. Due to the complex nature of their method, we have not been able to replicate it, although we hope to incorporate elements into future versions of our method.
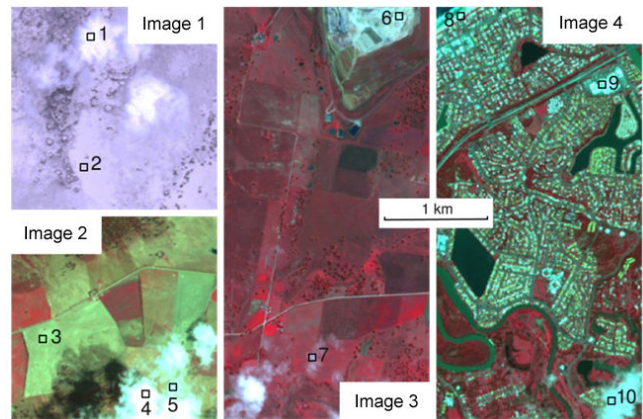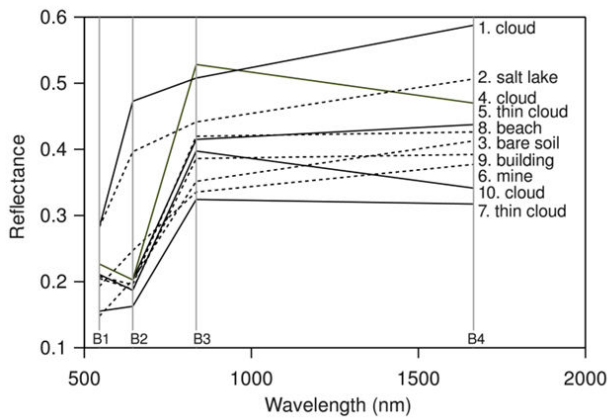
---

* Corresponding author.

Figure 2. Clouds and spectrally similar features in four SPOT-5 images from New South Wales, Australia. The graph shows bottom-of-atmosphere reflectance for pixels identified by the numbered black squares in the images. Cloud spectra are shown in solid lines, while other features are dashed lines. Images are shown with bands 3, 2 and 1 as RGB, and image numbers are those on Figure 1.

3. METHOD

The cloud and cloud shadow masking method presented here uses morphological feature extraction, where marker pixels are identified and mask segments are grown using the watershed from markers transformation. Manual input is only required to add and delete marker points in order to improve accuracy.

The method is split into two separate components, with details in the following sections. Firstly, a cloud mask is created, using the image as input. Secondly, the cloud shadow mask is created by using the image, acquisition parameters from the image metadata, and the cloud mask produced previously.

All image processing is implemented with the freely-available Python programming language, using functions from the standard Python library (Oliphant, 2007), as well as from the following open-source Python modules: Numpy (Oliphant, 2006), Scipy (Jones et al., 2001), Pymorph (Dougherty and Lotufu, 2003; Coelho, 2010), and GDAL (GDAL, 2009).

3.1 Clouds
Cloud masks are created using band 1 (green) of the SPOT5 data, as clouds appear to be consistently bright in this band compared to a variety of land cover types, and using a single band greatly simplifies the processing. Marker pixels for clouds are defined with a threshold in this band, which is determined by examining the band histogram. It is assumed that the right-most peak in the histogram (often near the maximum value) corresponds to cloud pixels, and so the position of the trough to the left of this peak is used as the threshold. This threshold often includes small features such as buildings, which are removed through converting the marker pixels to a binary image and subjecting it to morphological erosion. The position of the remaining marker pixels are then converted to a point vector file with a point at the central location of each pixel.

Cloud objects are grown from the marker pixels using the watershed from markers transformation of Felkel et al. (2001), shown in Figure 3. Dougherty and Lotufu (2003) present a detailed description of this morphological operation. First the morphological gradient of band 1 is calculated, giving an image where the edges of objects are high values and pixels within objects are low values. The watershed algorithm then grows segments out from the marker pixels until it reaches an edge with a neighbouring segment. This requires the neighbouring non-cloud pixels to also have marker pixels (referred to as external markers) so that the cloud segments are grown until they encounter an edge with non-cloud segments. After experimentation with different rules, we defined these external markers as all pixels greater than 900 m from a cloud marker. Before the segments are grown, connected groups of internal marker pixels are labelled with integers, while all external marker pixels are given the same integer label. This allows individual segments to be grown for each cloud, with a single large segment for the background.

Once the cloud segments are grown, they are converted to a binary image, smoothed and buffered. Smoothing is achieved through morphological opening, and is used to simplify any overly complex shapes not practical in an image mask. Buffering is achieved through morphological dilation to ensure that thin cloud pixels at the edges of clouds are included in the mask.

Generally, the method results in commission errors, such as large reflective buildings, concrete spillways below large dams, beach sand, playa lakes and salt crusts. It will also have omission errors, for thin clouds that appear semi-transparent. To avoid these errors, the marker point vector file can be edited, before re-running the watershed algorithm.
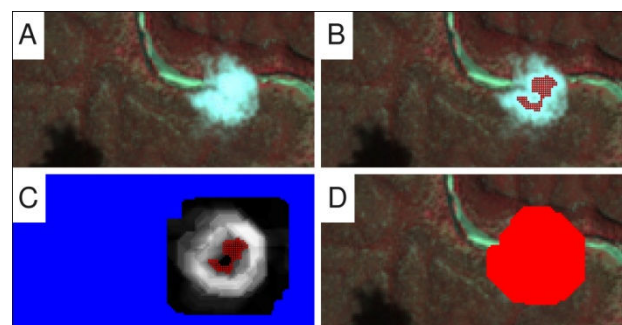

Figure 3. Stages in growing cloud segments using the watershed from markers transformation: (a) An image subset with bands 3, 2 and 1 as RGB; (b) Internal markers (red dots); (c) External markers in blue and the morphological gradient as greyscale; (d) Cloud segment in red.

3.2 Cloud-shadows

Cloud shadows are identified by searching for dark objects in the area likely to contain them. This requires the apparent solar azimuth (φ_a) to be calculated, using equation 1.

$$\phi_a = \pi + \arctan\left(\frac{\sin\phi_s \tan\theta_s - \sin\phi_v \tan\theta_v}{\cos\phi_s \tan\theta_s - \cos\phi_v \tan\theta_v}\right) \qquad (1)$$

where φ_s is the solar azimuth, θ_s is the solar zenith, φ_v is the viewer (sensor) azimuth, and θ_v is the viewer zenith. The area of likely shadow can then be determined by projecting the cloud mask along the apparent solar azimuth for a range of possible cloud heights. Complications due to variations in solar and viewing geometry across the scene, surface topography and the vertical shape of clouds are ignored. As cloud across the southern hemisphere's mid-latitudes is mostly low (Stubenrauch et al., 2005), cloud height is restricted to less than 3 km. This limits the area of likely shadow, simplifying the search for shadow pixels. If an image contains high cloud whose shadows are missed, the maximum height can be altered before re-running the algorithm.

The next stage identifies shadow marker pixels within the area of likely shadows. This, and all subsequent processes, uses band 3 (near-infrared) of the SPOT5 data, as cloud shadows appeared to be consistently dark in this band compared to a variety of land cover types. Marker pixels for cloud shadows are defined with a threshold in this band, which is determined by examining the band histogram. Firstly, a histogram is produced for the pixel values within the likely shadow area. This area is then buffered by 10 pixels to identify the surrounding area, which does not contain any clouds or shadows, and the histogram of pixel values from this area is produced. The subtraction of the second histogram from the first produces a graph of the pixel values most likely due to cloud shadow pixels. The peak of this graph is then identified, and used as the threshold for cloud shadow marker pixels. The position of the marker pixels are then converted to a point vector file to facilitate later editing if required.

External markers for cloud shadows are defined differently than the method used for clouds, as the simple distance criteria did not produce satisfactory results, with shadow segments commonly growing too large. To stop this, external markers are defined by local thresholds based on pixel values within 500 m from internal markers (where the local regions surrounding internal markers less than 500 m apart were combined). Within each region the external markers were defined as having a value greater than the local mean pixel value. This ensures that the shadow segments grow into locally dark pixels only. Once the internal and external markers are defined for the cloud shadows, segments were grown using the same method as described for clouds, including being smoothed and buffered.

4. RESULTS

The cloud and cloud shadow masking method has been applied to 142 cloudy SPOT5 images acquired during 2007 to 2009 over NSW (Figure 1). Processing was carried out on the NSW government high performance computing facility (4 machines, each with 2 Intel Xeon 3.00 GHz quad core processors and 64 GB of memory, running a 64 bit GNU/Linux operating system). Production of a mask for a single SPOT5 image (~55 million pixels) took around two minutes.

Editing time for the marker points was highly variable, ranging up to half an hour. Longer times were only necessary when an image contained many less reflective cloud objects. In these cases it was necessary to edit 2-3 times, re-growing the cloud segments between edits to identify problem areas.

Some examples of cloud and cloud-shadow masks are shown in Figure 4. The method is capable of masking most clouds and their shadows, dependant on the amount of cloud and the time available for editing marker pixels. It cannot generally mask very transparent clouds, and has difficulty with cloud and shadow segments that overly features with similar reflectance properties.
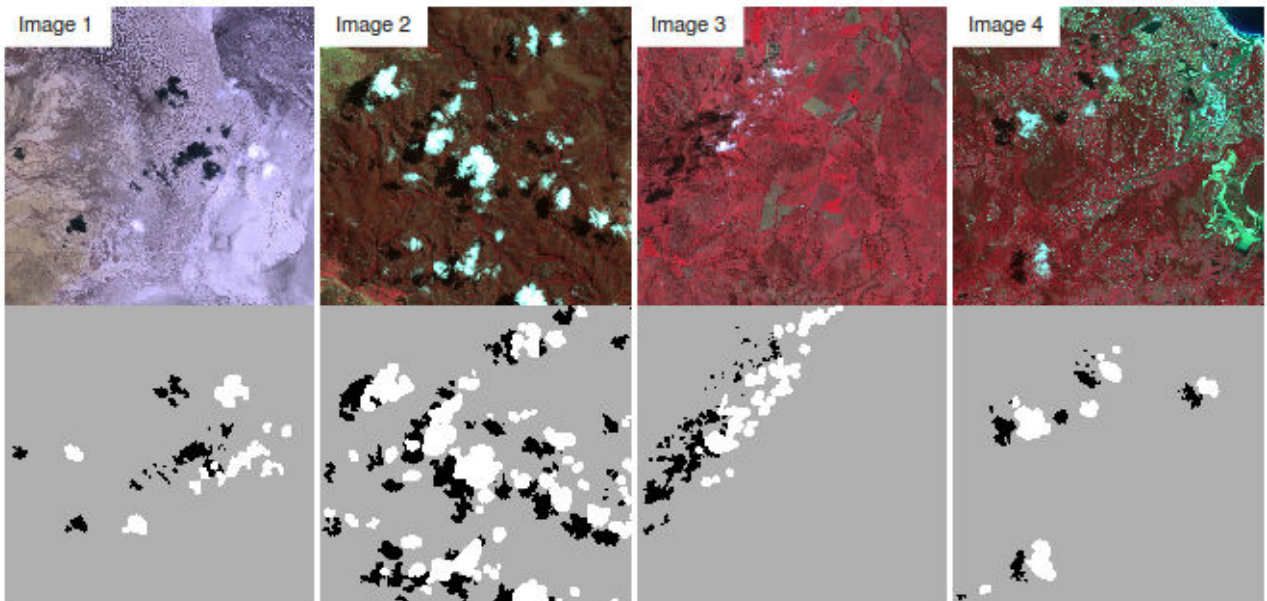
Figure 4. Subsets of SPOT5 images (top) and their corresponding cloud and cloud shadow masks (below). Subsets are 10 km by 10 km, shown with bands 3, 2 and 1 as RGB, while masks are coloured with clouds as white and cloud shadows as black.

Table 1. Overall accuracies of the cloud and cloud shadow masking method, and the commission and omission errors for the cloud and shadow classes. All accuracy and error values are percentages (%), and image numbers refer to Figure 1.

Image	Cloud		Shadow		Overall accuracy
	Commission	Omission	Commission	Omission	
1	26	1	6	0	89
2	23	1	22	9	83
3	46	4	18	6	78
4	26	4	34	0	79
Total	30	2	20	4	82

As alternative data on the distribution of clouds is not available, validation of the method against independent data was not possible. Instead, a set of 300 manually classified reference pixels were sampled from each image. Sampling was stratified equally between pixels classified as cloud, shadow, and those within 5 km of a cloud/shadow, to ensure that the assessment targeted cloudy areas. The results confirm the observation that although omission errors are low, ranging from 1-9 %, the commission errors are larger, ranging from 6-46 % (Table 1). Differences in mask accuracy between images relates mostly to the number of cloud objects (e.g. image 3 has more clouds). While the overall accuracy is reasonably high, at approximately 80 %, future improvements are anticipated, and are discussed below.

5. FUTURE IMPROVEMENTS

The cloud and cloud shadow masking method presented here has proven to be relatively fast and robust, working on images that have a large variety of land cover and cloud types. It has also been successfully applied to 29 Landsat TM/ETM+ images, facilitating the calculation of FPC and the detection of changes to the extent of woody vegetation across NSW.

The method was designed for immediate application, and has several known limitations. Some areas being investigated for future improvement are incorporating multiple image bands; using image subsets to search for local thresholds, and; the pairing of cloud and shadow objects. For example, 2-dimensional density plots are being used to establish consistent band-ratio thresholds for defining external marker pixels, which will improve segmentation, smoothing and buffering.

REFERENCES

L.P. Coelho, "Pymorph: Python Morphology Toolbox," 2010. http://luispedro.org/software/pymorph

T. Danaher, P. Scath, J. Armston, L. Collett, J. Kitchen, and S. Gillingham, "Remote Sensing of Tree-Grass Systems: the Eastern Australian Woodlands," in M.J. Hill, and N.P. Hanan (Eds.), Ecosystem Function in Savannas: Measurement and Modeling at Landscape to Global Scales, p.p. 175-194, CRC Press, Boca Raton, 2010.

E.R. Dougherty and R.A. Lotufu, "Hands-on morphological image processing," SPIE Press, Washington, 2003.

P. Felkel, et al., "Implementation and complexity of the watershed-from-markers algorithm computed as a minimal cost forest," Computer Graphics Forum, vol. 20, 2001.

GDAL, "GDAL - Geospatial Data Abstraction Library," Open Source Geospatial Foundation, 2009. http://www.gdal.org/

C. Huang, N. Thomas, S.N. Goward, J.G. Masek, Z. Zhu, J.R.G. Townshend, and J.E. Vogelmann, "Automated masking of cloud and cloud shadow for forest change analysis using Landsat images," International Journal of Remote Sensing, vol 31, p.p. 5449-5464, October 2010.

E. Jones, et al., "SciPy: Open Source Scientific Tools for Python," 2001. http://www.scipy.org/

S. Le Hégarat-Mascle, and C. André, "Use of Markov Random Fields for automatic cloud/shadow detection on high resolution optical images," ISPRS Journal of Photogrammetry and Remote Sensing, vol 64, p.p. 351-366, July 2009.

A. Meier, T. Danaher, T. Gill, S. Smith, G. Horn, and N. Flood, "State-wide inter-annual changes to foliage projective cover: better products from higher resolution satellite imagery and improved processing methodologies," 34th International Symposium on Remote Sensing of Environment, Sydney, Australia, April 2011.

T.E. Oliphant, "Guide to NumPy," Trelgol Publishing, 2006. http://www.tramy.us/numpybook.pdf

T.E. Oliphant, "Python for Scientific Computing," Computing in Science & Engineering, vol. 9, p.p. 10-20, 2007.

C.J. Stubenrauch, F. Eddounia, and L. Sauvage, "Cloud heights from TOVS Path-B: Evaluation using LITE observations and distributions of highest cloud layers," Journal of Geophyscial Research, vol 110, D19203, doi:10.1029/2004JD005447, 2005.

ACKNOWLEDGEMENTS