Rapid detection of new and expanding human settlements in the Limpopo province of South Africa using a spatio-temporal change detection method

W. Kleynhans[a,c], B.P. Salmon[c,d], K.J. Wessels[c,b], J.C. Olivier[d]

[a] *Department of Electrical, Electronic and Computer Engineering, University of Pretoria, South Africa.*
[b] *Department of Geography, University of Pretoria, South Africa.*
[c] *Remote Sensing Research Unit, Meraka Institute, CSIR, South Africa.*
[d] *School of Engineering, University of Tasmania, Australia*

Abstract

Recent development has identified the benefits of using hyper-temporal satellite time series data for land cover change detection and classification in South Africa. In particular, the monitoring of human settlement expansion in the Limpopo province is of relevance as it is the one of the most pervasive forms of land-cover change in this province which covers an area of roughly 125 000km^2. In this paper, a spatio-temporal autocorrelation change detection (STACD) method is developed to improve the performance of a pixel based temporal Autocorrelation change detection (TACD) method previously proposed. The objective is to apply the algorithm to large areas to detect the conversion of natural vegetation to settlement which is then validated by an operator using additional data (such as high resolution imagery). Importantly, as the objective of the method is to indicate areas of potential change to operators for further analysis, a low false alarm rate is required while achieving an acceptable probability of detection. Results indicate that detection accuracies of 70% of new settlement instances are achievable at a false alarm rate of less than 1% with the STACD method, an improvement of up to 17% compared to the original TACD formulation.

Keywords: change detection, Autocorrelation, time-series, hyper-temporal, settlements

Email address: `wkleynhans@csir.co.za` (W. Kleynhans)

1. Introduction

The most pervasive form of land-cover change in South Africa and many other developing countries around the world is human settlement expansion. In many cases, new human settlements as well as existing settlements expand informally and these expansions occur in areas that were previously covered by natural vegetation. Informal or unplanned settlements usually are formed as people move closer to employment opportunities, or in response to environmental and/or market forces. These settlements occur in various locations and often lack basic services such as electricity, running-water, water-borne sewage and refuse removal. The spatial layout is often unplanned and informally developed by the inhabitants of the settlements themselves [1]. Figure 1 shows an informal settlement in the Limpopo province of South Africa which developed between the years 2003 and 2009 in an area that was initially mostly covered by natural vegetation. The development of these settlements need to be detected so that they can be mapped in detail and accommodated during planning sessions undertaken by regional and local government.

Detailed mapping of settlements are usually done by analysts digitizing features from aerial photography or high resolution satellite images at 1:10 000 scale. The feature data sets need to be updated regularly (at least every two years) to support spatial planning, especially where it concerns new settlements. Updating maps over large areas using manual photo-interpretation to scan the entire area is slow and costly. Due to constraints in cost and resources experienced by most mapping agencies, feature datasets can be up to a decade old, while only a small percentage of the area actually experienced change. Methods that can rapidly indicate areas having a high probability of change is thus very valuable to analyst as this can be used to direct their attention to high probability change areas for further evaluation using, for example, higher resolution imagery of the area. By using such a targeted approach, an increase in mapping efficiency of up to ten times has been observed compared to a complete re-extraction [2]. We therefore focus on the development of automated change detection methods based on hyper-temporal satellite imagery to ultimately try to improve the productivity of detailed mapping efforts.

Satellite time series data has proven to be an effective data source for change detection [3, 4, 5, 6] and in particular, time series analyses of hyper-temporal satellite data has been successfully applied for land cover change detection in South Africa specifically related to the monitoring of human settlement expansion. In [7], a Neural network based post classification change detection approach was used to detect when land cover conversion takes place from natural vegetation to settlement classes. In [8], MODIS time-series data was modeled as a triply modulated cosine function and an Extended Kalman filter was used to track the parameters of the model and declare change based on parameter behavior. In [9], the use of Page's cumulative sum (CUSUM) test was proposed as a method to detect new settlement. It should be noted that these aforementioned methods make limited use of spatial information and are predominantly pixel based.

An autocorrelation function (ACF) change detection method was recently demonstrated to detect the development of new human settlements in South Africa [10]. This method uses MODIS time-series data, which have previously been shown to be separable (distinguishable) for the natural vegetation and settlement land cover classes considered in this study [11]. The method uses the ACF of a MODIS time-series to provide an indication of the level of time-series stationarity (by considering the stability of the time-series mean and variance over time) which is then consequently used as a measure of land cover change.

In the original formulation of the ACF approach [10], a single pixel's entire time-series for a single band (spanning eight years) was used as input. A change metric was then calculated by considering the properties of the ACF of the time-series. When the resulting change index was compared to a threshold value, a per-pixel based change alarm resulted. In this paper, the aforementioned method is extended to incorporate spatial information by not only considering the change index for a single pixel but also that of its surrounding pixels to determine whether a change should be declared. The proposed Spatio Temporal ACF change detection (STACD) method uses temporal only ACF change metrics as calculated using the approach in [10] on a per pixel basis and compares this temporal ACF change index with that of pixels in its neighborhood to increase performance. It is important to note that, the method is adapted to be able to easily incorporate multiple MODIS bands. The idea behind the Spatio-Temporal extension to the classic ACF change detection method is that when a single threshold is used over a large area (as was done in [10]) there will inevitably be areas that are more non-stationary due to, for example, drought in arid areas, large scale commercial agriculture, etc. This will result in large areas showing up as change in regional maps that are not necessarily the type of anthropogenic changes that are of interest. The spatio-temporal extension to the ACF change detection method mitigate these changes by considering the change properties of the pixel neighborhood.

The objective of this paper is to develop a robust change detection method that is able to detect the formation of new informal settlements in areas that are typically covered by diverse natural vegetation. The detected changes should then be used to alert operators to areas of possible changes which could thereafter be validated, and the necessary maps updated, using high resolution imagery. Although the false alarm rate of the method can be set to any percentage depending on the requirement of the operator, in our use case the false alarm rate requirement was very low ($< 1\%$) as the area on which the change algorithm is run is large and the validation of a large number of false alarms could be very costly and time consuming.

2. Data Description

2.1. MODIS data

The time series data used in this study was derived from MODIS data. The MODIS instrument data are converted systematically into terrestrial, atmospheric and oceanic products. The first seven bands (covering spectral bands in the visible, near infrared and shortwave infrared range [12]) are typically used for land applications and are often referred to as the MODIS land bands. The specific land product that was used was the MCD43A4 product [13]. This specific product have been used in various land cover change detection applications [7, 8, 9] and is produced using data acquired from the MODIS sensor on-board the Aqua and Terra satellites and provides one composited sample (consisting of 16 days of acquisition) every 8 days. This product has a spatial resolution of 500m and is BRDF-corrected. The time period that was considered is 2001/01 to 2008/01 resulting in a time-series containing 315 MODIS observations. Quality Control (QC) flags where not explicitly used in the preprocessing of the time series data but it should be noted that the study area that was considered does not have prolonged time-periods of cloud cover which results in the data not containing significant missing values (typically less than 4% of samples). In the rare occurrence of missing values, cubic spline interpolation was used to infer these missing values [14].

2.2. Study Area

The study area is located in northern South Africa and is mostly covered by natural vegetation which predominantly consist of grassland, savanna and shrub-land. A large number of informal settlements are however rapidly expanding throughout the area. The study area covers an approximate 25000 km^2 having an upper left coordinate of (23°20′12.09″S ; 28°35′25.18″E) and a lower right coordinate of (25°00′14.59″S ; 30°06′58.30″E). A total of 1497 examples of natural vegetation 500m MODIS pixels were identified within the study area. Each of these pixels were evaluated using SPOT5 high resolution data to ensure that none of them have experienced any land cover change during the study period. Examples of confirmed settlement developments (i.e. change from natural vegetation to settlement) were obtained by means of visual interpretation of high resolution Landsat and SPOT images in 2000 and 2008 respectively. All settlements identified in 2008 were referenced back to the same physical area in 2000 and all the new settlement polygons were mapped. A total of 30 occurrences of settlement development were identified and the corresponding MODIS pixels (n=*117*) were so identified. At least 70% of the total area within a pixel had to have changed for inclusion into the change dataset. Figure 2 shows the location of the Limpopo province as well as the ground truth pixels used in the study.

3. Methodology

3.1. Temporal ACF Change detection (TACD) method

The Temporal ACF change detection (TACD) method proposed in [10] uses a two stage approach. Firstly, the band, lag and threshold selection is done using a simulated change dataset together with a no-change dataset. Second, the aforementioned parameters are used in an unsupervised manner to detect change. Assume that a MODIS time-series is expressed as

$$\mathbf{X} = \{X_n\}_{n=1}^{n=T}, \tag{1}$$

where X_n is the observation from an arbitrary spectral band at time n and T is the number of time-series observations available. The ACF for time-series $\mathbf{X}$ can then be expressed as

$$R(\tau) = \frac{E[(X_n - \mu)(X_{n+\tau} - \mu)]}{\mathrm{var}(\mathbf{X})}, \tag{2}$$

where τ is the time-lag and E denotes the expectation. The mean of $\mathbf{X}$ is given as μ and the variance, which is used for normalization, is defined as $\mathrm{var}(\mathbf{X})$. The mean and variance of the time-series of $\mathbf{X}$ in (2) is required to remain constant through time to determine the true ACF of the time-series. The inconsistency of the mean and variance typically associated with a change pixel's non-stationary time-series becomes apparent when analyzing the ACF of the time-series. An example of this is shown in figure 3, where MODIS band 1, 3 and 4 time-series is shown for a change and no-change pixel together with their corresponding ACF. The change metric was defined in [10] as the temporal correlation at a specific lag (τ) given as

$$R(\tau) = \delta_\tau. \tag{3}$$

Using a change threshold (δ_τ^*), a change or no-change decision was made as

$$\text{Change} = \begin{cases} \text{true} & \text{if } R(\tau) \geq \delta_\tau^* \\ \text{false} & \text{if } R(\tau) < \delta_\tau^* \end{cases}$$

The value of τ as well the threshold value (δ_τ^*) was determined by using simulated change and no-change datasets after which the resulting parameters were used to run the algorithm in an unsupervised manner for the entire study area [10].

3.2. Spatio Temporal ACF change detection (STACD) method

An important note on the TACD method is that although any band can be used by the method, the method is not inherently adaptable to multi-band data. Second, the method only uses a single time series (using one of the MODIS spectral bands) for each pixel as input and outputs a change / no-change decision, thus only using information related to a single pixel as input and not utilizing any spatial information. Some natural vegetation types are more stationary than others on intra-annual / seasonal and inter-annual basis. The diverse regional distribution of this background stationarity may thus result in false alarms if a single threshold was applied to the change metric for an entire region on a per-pixel basis and no spatial context is used. Third, the TACD method requires the use of simulated change to determine the autocorrelation lag (τ value in (3)), which is an additional parameter that needs to be estimated and could potentially have a negative effect on the accuracy if not estimated correctly.

In this section, an extension to the TACD is proposed to address the aforementioned shortcomings. First, a modification to the change index value (as given in (3)) is made as

$$\delta^b_{x,y} = \sum_{\tau=1}^{k} R(\tau, x, y, b) = \sum_{\tau=1}^{k} \frac{E[(X_n^{x,y,b} - \mu)(X_{n+\tau}^{x,y,b} - \mu)]}{\mathrm{var}(\mathbf{X})}, \quad (4)$$

Where b is the MODIS band, $X^{x,y,b}$ is the time-series of band b at pixel location (x,y), τ is the ACF lag and k is the total number of lags to be summed. From (4) it is evident that the new change metric uses a summation of the first k lags of $R(\tau)$. By using the summation, the method is not as sensitive to the selection of a specific value of τ. The specific value of k was calculated by sweeping all possible values of k and evaluating the overall accuracy. It was found that adding more than 23 lags did not increase the overall accuracy. Consequently a value of $k = 23$ was used in this study. It was found that there was only a marginal reduction in accuracy compared to the case where the value of τ is explicitly specified. The k value was fixed and the same value was used throughout the study. This implied that the specific value of τ no longer needed to be estimated, thus making the method more general. A new change index which incorporates both spatial as well as multiple spectral bands was then formulated as follows:

Consider a neighborhood of pixels around a pixel located at position (x, y) with radius N as

$$X^b_{N,x,y} = \begin{bmatrix} \delta^{(b)}_{(x-N)(y-N)} & \cdots & \delta^{(b)}_{(x+N)(y-N)} \\ \vdots & \delta^{(b)}_{xy} & \vdots \\ \delta^{(b)}_{(x-N)(y+N)} & \cdots & \delta^{(b)}_{(x+N)(y+N)} \end{bmatrix}, \quad (5)$$

where $\delta^{(b)}_{xy}$ is the change index as calculated in (4) for band b pixel (x, y). A new change metric ($\gamma_{x,y}$) is proposed that first calculates the mean value of $X^b_{N,x,y}$ (excluding the center pixel) as

$$s_{x,y,b} = \frac{\sum X^b_{N,x,y} - \delta^{(b)}_{xy}}{(2N+1)^2 - 1}, \tag{6}$$

and then calculates the Euclidean distance between the center pixel and mean value

$$\gamma_{x,y} = \sqrt{\sum_b (\delta^{(b)}_{xy} - s_{x,y,b})^2}. \tag{7}$$

The new change metric ($\gamma_{x,y}$) is thus a Euclidean distance in multidimensional space between the mean change metric for all the pixels in the neighborhood and the center pixel. Thus, when a neighborhood of pixels are inherently non-stationary, the average change index for all pixels in the neighborhood will subsequently also be high and the calculated Euclidean distance change index will in effect be normalized with reference to the neighborhood. This spatial adaptation allows locally adaptive scaling so that settlement expansion (or other anthropogenic land cover transformations), are highlighted relative to the natural change the larger region may be experiencing. This results in a much more robust framework with a reduced false alarm rate compared to the original formulation that did not consider spatial information. It follows that, when comparing the change metric ($\gamma_{x,y}$) with a threshold value, large areas of change that are not typical for the change relevant in this study, will have a reduced Euclidean distance and therefore a smaller chance of being declared as change. Figure 4 shows a scatter plot and distribution heat-map of $X^b_{N,x,y}$ for $b = \{3, 4\}$ and $N = 15$. The center pixel δ_{xy} corresponds to the change pixel time-series shown in figure 3. It can be seen that there is a significant distance between δ_{xy} and the mean of the majority of the change index values.

4. Results

The evaluation that was done on the results of the method was done in relation to the Change detection accuracy (CDA) (defined in this study as the percentage of known change pixels detected as change) as well as the false alarm rate (defined in this study as the percentage of known unchanged pixels detected as change). The accuracy assessment was done using ground

Table 1: Change Detection Accuracy (CDA), False Alarm Rate (FAR) and overall accuracy (O_A) results for the temporal ACF Change detection (TACD) algorithm given for each of the MODIS bands. The number of pixels (n) used in the change dataset as well as the no-change dataset is also given in the CDA and FAR column respectively.

Band	CDA ($n = 117$)	FAR ($n = 1497$)	O_A
1	34%	1%	67%
2	6%	1%	52%
3	33%	1%	66%
4	33%	1%	66%
5	4%	1%	52%
6	9%	1%	54%
7	13%	1%	56%

Table 2: Spatio-Temporal ACF Change detection (STACD) algorithm's Change Detection Accuracy (CDA) ($n = 117$) as a function of the value of neighborhood size ($N = \{3, 5, 10, 15, 20\}$) at a False Alarm rate of 1% ($n = 1497$). Overall accuracy in each instance is given in brackets.

Band	N=3	N=5	N=10	N=15	N=20
1	44% (72%)	50% (75%)	**51**% (**75**%)	47% (73%)	44% (72%)
2	9% (54%)	9% (54%)	9% (54%)	9% (54%)	9% (54%)
3	41% (70%)	43% (71%)	39% (69%)	32% (66%)	29% (64%)
4	39% (69%)	39% (69%)	42% (71%)	39% (69%)	36% (68%)
5	14% (57%)	13% (56%)	14% (57%)	12% (56%)	12% (56%)
6	18% (59%)	19% (59%)	15% (57%)	16% (58%)	17% (58%)
7	31% (65%)	29% (64%)	28% (64%)	23% (61%)	21% (60%)

truth dataset with a spatial distribution over the Limpopo province shown in figure 2. As the requirement is to have a False alarm rate less than 1%, the change detection accuracy at a 1% False alarm rate was considered. The results for both the TACD method as well as the STACD method is presented in tables 1 and 2. In table 2 the neighborhood size parameter (N) was varied between 3 and 20. It is clear that the inclusion of spatial information has improved the CDA for each of the MODIS bands regardless of the value of N. The maximum CDA was achieved using band 1 with a value of $N = 10$. Multiple band combinations as well as vegetation indices including NDVI and EVI where also considered, but it was found that no significant improvement was achievable over that of using only band 1. Band 1 is the Red band in the visible spectrum range 620-670nm and is known to be very sensitive to changes in vegetation. This is not to say that band 1 will always perform better, but only that for our specific land cover conversion case (natural vegetation being converted to human settlement) and study area, using only band 1 gave the best results. As land cover heterogeneity and land cover type and change dynamics will vary significantly between applications, it is advised that different values of N as well as various band combinations be considered when applying the algorithm to different scenarios. It is clear that the STACD algorithm has a better performance for the Limpopo dataset with an improved overall CDA accuracy of nearly 17% compared to the best result using the original TACD formulation when considering a false alarm rate of $< 1\%$. This also compared well to NDVI differencing [4] which achieved a 32% CDA at the same false alarm rate in this study area. At first glance a 51% change detection accuracy could be regarded as being relatively low but it should be noted that this figure should be seen in context of the the low false alarm rate (1%) (overall accuracy of 75%), which is nearly a 10% improvement compared to the original method (table 1).

In addition, new informal settlements very often affect a group of contiguous pixels and even if only one of these pixels are declared as change, an operator would be alerted and the surrounding area would be analyzed using high resolution imagery. Figure 5 shows the distribution of the number of MODIS pixel contained in each of the 30 settlement polygons identified. It was found that more than 73% of new or expanding settlements affected more than one MODIS pixel, thus increasing the probability of detection of the new settlement instance. To illustrate this point, the combination giving the best result (as shown in table 2) with $b = 1$ and $N = 10$ was applied over the entire study area, the pixels identified as change by the method was compared to the confirmed settlement polygons in the test dataset. It was found that 70% of the settlements in the test dataset were found to have at least one MODIS pixel indicated as having changed and was correctly identified by the method as a new or expanding settlement instance (1% false alarm rate and overall accuracy of 85%). Figure 6 shows an example of an informal settlement expansion with a size of approximately 6 MODIS pixels that occurred in the Limpopo province between 2005 (top image) and 2008 (bottom image). The known change polygon (shown in blue), is the total extent of the change whereas the detected change (shown in red) is the pixels detected as change by the algorithm. Even though three of the pixels (50%) were not flagged by the method as having changed, it is clear to an operator investigating this scene that a change has occurred and that this settlement expansion instance has successfully been identified.

It should also be noted that, traditionally, spatial temporal methods can be extremely time consuming and do not scale particularly well when extended regionally, for this reason the method propose in this paper used a computationally efficient two step approach. Firstly the ACF change index is calculated for all pixels in the area of interest, this involves simply computing the ACF of each time series and computing the sum over the first k lags (equation)). The second step uses a sliding window (based on the window size N) over each of the pixels and calculating the mean of the neighborhood and euclidean distance from the pixel in question to that calculated mean. The low computational complexity of steps one and two allows the method to be very scalable and allows the method to be applied to large areas.

5. Conclusion

In this paper, an extension is proposed to the temporal autocorrelation change detection (TACD)method proposed in [10]. The original formulation is a threshold based change alarm that uses an ACF of a time-series to infer a change index. The method proposed in this paper is an extension to the aforementioned which utilizes spatial context to detect the formation of new informal settlements in areas that are typically covered by diverse natural vegetation. This change detection framework is intended to be used as a tool to alert operators to areas of possible changes between two dates which could then be validated using a secondary step (such as the use of high resolution imagery). As the algorithm is intended to be run over potentially very large areas (regional scale), a primary objective was to ensure that a very low false alarm rate should be maintained. It was shown that the addition of spatial information enabled a much lower false alarm rate ($<$ 1%) to be achieved while detecting 70% of the new or expanding settlements in the study area. As mentioned previously, one of our major considerations was to have a very low false alarm rate, as the requirement of a low false alarm rate is critical in the manual validation of rapidly detected change areas (i.e when running the method over large areas containing hundreds of thousands of pixels, the difference between a false alarm rate of only a few percent could result in hundreds of false alarms that would then have to be manually validated by an operator which could take a significant amount of time). By having a constraint on the false alarm that can be tolerated by an operator (in our case 1%), we calculated the corresponding threshold at a 1% false alarm rate using a no-change dataset. The implication of this in an operational environment is that when selecting a threshold value, only the false alarm rate requirement needs to be specified together with an example dataset of no-change examples only which are typically easy to obtain as most areas do not change significantly over time. Although specific change date information is not provided, this algorithm could still be used effectively to rapidly determine areas of high change probability between two dates. An example would be mapping agencies wanting to update maps between mapping intervals. Although the spatio-temporal autocorrelation change detection (STACD) algorithm has only been tested for the case of new or expanding settlement detection the Limpopo province (located in the most northern part of South Africa), the method can easily be applied regionally and has been applied across the entire South Africa and is in the process of being accessed by the official mapping agency (NGI) using aerial maps. Results of this national validation will be reported in future publications

[1] A. Palframan, A syntactical analysis of settlement form an investigation of socio-spatial characteristics in low-income housing settlements in Port Elizabeth, South Africa., in: World Congress on Housing, Vol. 33, Pretoria, South Africa, 2005, pp. 3353–3355.

[2] R. Mitchell, Change-driven mapping, Geospatial Intelligence Forum 10 (4) (2012) 22–24.

[3] J. Verbesselt, R. Hyndman, G. Newnham, D. Culvenor, Detecting trend and seasonal changes in satellite image time series, Remote Sensing of Environment 114 (1) (2010) 106–115.

[4] R. S. Lunetta, J. F. Knight, J. Ediriwickrema, J. G. Lyon, L. D. Worthy, Land-cover change detection using multi-temporal MODIS NDVI data, Remote Sensing of Environment 105 (2) (2006) 142–154.

[5] K. de Beurs, G. Henerby, A statistical framework for the analysis of long image time series, International Journal of Remote Sensing 26 (8) (2005) 1551–1573.

[6] R. Kennedy, W. Cohen, T. Schroeder, Trajectory-based change detection for automated characterization of forest disturbance dynamics, Remote Sensing of Environment 110 (3) (2007) 370–386.

[7] B. Salmon, J. Olivier, W. Kleynhans, K. Wessels, F. van den Bergh, K. Steenkamp, The use of a multilayer perceptron for detecting new human settlements from a time series of modis images, International Journal of Applied Earth Observation and Geoinformation 13 (6) (2011) 873–883.

[8] W. Kleynhans, J. C. Olivier, K. J. Wessels, B. P. Salmon, F. van den Bergh, K. Steenkamp, Detecting land cover change using an extended kalman filter on MODIS NDVI time-series data, IEEE Geoscience and Remote Sensing Letters 8 (3) (2011) 507–511.

[9] T. Grobler, E. Ackermann, A. van Zyl, J. Olivier, W. Kleynhans, B. Salmon, Using page's cumulative sum test on modis time series to detect land-cover changes, IEEE Geoscience and Remote Sensing Letters 10 (2) (2013) 332–336.

[10] W. Kleynhans, B. Salmon, J. Olivier, F. van den Bergh, K. Wessels, T. Grobler, K. Steenkamp, Land cover change detection using autocorrelation analysis on modis time-series data: Detection of new human settlements in the gauteng province of south africa, Selected Topics in Applied Earth Observations and Remote Sensing, IEEE Journal of 5 (3) (2012) 777 –783. doi:10.1109/JSTARS.2012.2187177.

[11] W. Kleynhans, J. C. Olivier, K. J. Wessels, B. P. Salmon, F. van den Bergh, K. Steenkamp, Improving land cover class separation using an extended Kalman filter on MODIS NDVI time series data, IEEE Geoscience and Remote Sensing Letters 7 (2) (2010) 381–385.

[12] E. Vermote, N. E. Saleous, C. Justice, Atmospheric correction of MODIS data in the visible to middle infrared: first results, Remote Sensing of Environment 83 (1-2) (2002) 97 111.

[13] C. Schaaf, et al., First operational BRDF, albedo nadir reflectance products from MODIS, Remote Sensing of Environment 83 (1-2) (2002) 135–148.

[14] K. Atkinson, On the order of convergence of natural cubic spline interpolation, SIAM Journal on Numerical Analysis 5 (1) (1968) 89–101.

Figure 1: Example of a new settlement development in the Limpopo province of South Africa. QuickBird image on the left shows the area being mostly covered by natural vegetation in 2003, whereas the Quickbird image on the right shows a new informal settlement in 2009 (courtesy of Google$^{\text{TM}}$Earth). The red polygons show the footprint of three 500m $\times$ 500m MODIS pixels.

Figure 2: Provincial map of South Africa showing the location of the the Limpopo province (left). Zoom in of the Limpopo province showing the ground truth pixels used in this study are shown on the right. The area corresponding to the example shown in figure 1 is indicated with the blue circle.

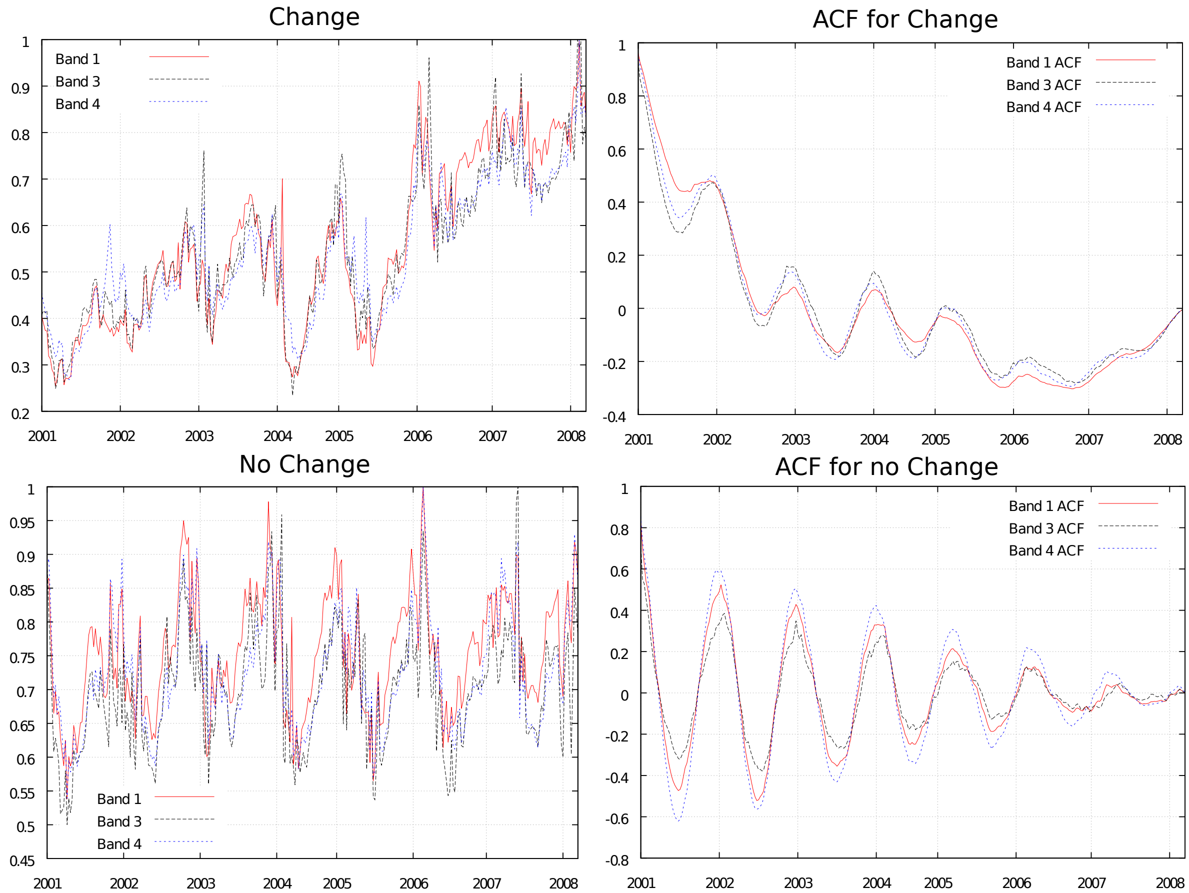

Figure 3: MODIS band 1,3 and 4 time-series for a pixel that undergone change (top-left) and no-change (bottom left) respectively. The ACF of the change time-series (top right) and no-change time-series (bottom right) is also shown.

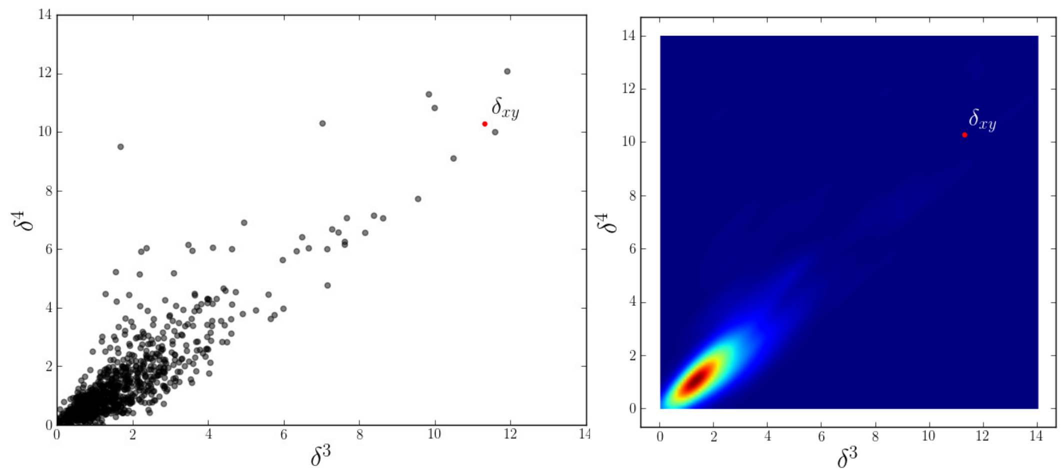

Figure 4: Example of the change index values for band 3 (δ^3) vs. band 4 (δ^4) for a neighborhood of $N = 15$ pixels (total of 225 pixels) around the center pixel. This example shows a center pixel (x, y) undergoing a high degree of change corresponding to the change time-series shown in Figure 3 .The image on the left shows a scatter plot of δ^3 vs. δ^4 and the image on the right shows the corresponding distribution heat-map.

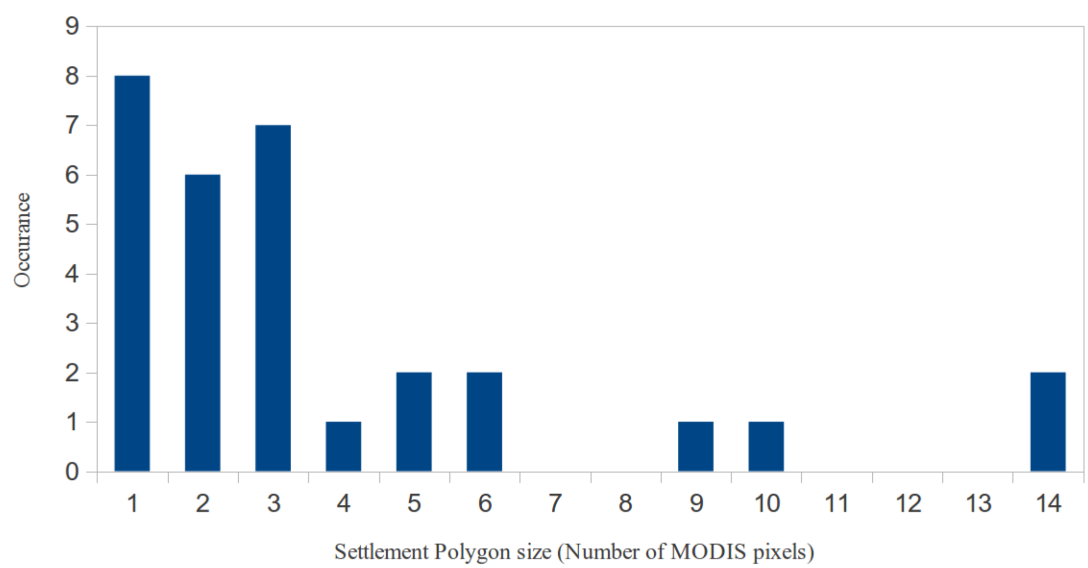

Figure 5: Distribution of the settlement sizes (given as the number of MODIS pixels) of the test dataset.

Figure 6: Example of informal settlement expansion in the Limpopo province between 2005 (top image) and 2008 (bottom image). Ground truth information indicate that 6 MODIS pixels (pixels 1-6) were affected by this change. The blue polygons (pixels 1,2 and 6) were not detected as change by the proposed detection algorithm where the red polygons (pixels 3,4 and 5) were detected as having changed. It can be seen that even though only half of the change pixels were detected (i.e detection accuracy of 50%), this new settlement instance was detected and could consequently be mapped with high precision by an operator using high resolution imagery over the area (courtesy of Google$^{\mathrm{TM}}$Earth).