A MODEL-FREE TIME SERIES SEGMENTATION APPROACH FOR LAND COVER CHANGE DETECTION

ASHISH GARG†*, LYDIA MANIKONDA†*, SHASHANK KUMAR**, VIKRANT KRISHNA*, SHYAM BORIAH*, MICHAEL STEINBACH*, VIPIN KUMAR*, DURGA TOSHNIWAL**, CHRISTOPHER POTTER***, AND STEVEN KLOOSTER***

Abstract. Ecosystem-related observations from remote sensors on satellites offer significant possibility for understanding the location and extent of global land cover change. In this paper, we focus on time series segmentation techniques in the context of land cover change detection. We propose a model-based time series segmentation algorithm inspired by an event detection framework proposed in the field of statistics. We also present a novel model-free change detection algorithm for detecting land cover change that is computationally simple, efficient, non-parametric and takes into account the inherent variability present in the remote sensing data. A key advantage of this method is that it can be applied globally for a variety of vegetation without having to identify the right model for specific vegetation types. We evaluate the change detection capacity of the proposed techniques on both synthetic and MODIS EVI data sets. We illustrate the importance and relative ability of different algorithms to account for the natural variation in the EVI data set.

1. Introduction

The goal of the land cover change detection problem is to detect when the land cover at a given location has been converted from one type to another. It is very important to study land cover change in order to understand its impact on local climate, radiation balance, biogeochemistry, hydrology, and the diversity and abundance of terrestrial species [6, 17]. Such understanding can be very valuable for policy makers, natural resource managers and researchers to address the issues related to global environmental changes. A large body of change detection studies from remotely sensed imagery has focused on comparisons between two images: one before and one after a change [8]. However, such techniques are usually domain or region specific and require expensive training and thus are difficult to scale globally. Recognizing these limitations, several algorithms have been developed [17, 16] to detect changes in the time series of satellite-based observations such as the Enhanced Vegetation Index (EVI) [3]. EVI, which is a product based on measurements taken from the MODIS instrument on NASA's Terra and Aqua satellites, is available globally at 250m and 1km resolution and at a temporal frequency of 16 days, since February 2000.

A number of techniques [5, 11] have been developed recently for identifying sudden drops in the vegetation index time series (e.g. in Figure 1(a)) or slow degradation (e.g. in Figure 1(b)) that can occur due to fires or logging etc. However, these techniques are unable to effectively detect changes such as conversion of forested land to crop land, intensification of agriculture, and change in cropping patterns. These changes do not necessarily result in loss of vegetation, for example, see Figure 1(c) for the change in cropping pattern from double to single crop per year. Rather, these changes result in characteristic change in the regular pattern of the EVI time series. The ability to monitor such land cover changes at local, regional and global scale is important due to their potential impact on the environment.

† These authors contributed equally to this work.

*University of Minnesota, <ashish,lmani,vikrant,sboriah,steinbac,kumar>@cs.umn.edu
**Indian Institute of Technology, Roorkee <sha28uec,durgafec>@iitr.ernet.in
***NASA Ames Research Center, chris.potter@nasa.gov, sklooster@gaia.arc.nasa.gov .

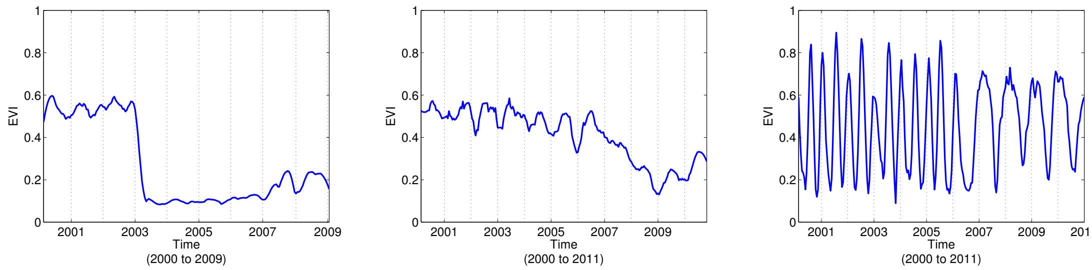

(a) (b) (c)

FIGURE 1. EVI vegetation time series (Feb–2000 to Sep–2010 -Vertical lines indicate yearly boundaries) showing (a) Sudden drop in year 2003 at a location in California; (b) Slow degradation from 2006 to 2009 at a location in Washington; (c) Conversion of double to single cropping for a location in Zimbabwe in 2006. Note that the mean EVI for each year is similar for this time series.

The problem of detecting land cover changes can be posed as segmenting a vegetation index time series. The goal of segmentation is to partition the input time series into homogeneous segments such that the subsequence within a segment is homogeneous and the segments are heterogenous with respect to each other. Segmentation thus is essentially a special case of change detection since by definition successive segments are not homogeneous.

In this paper, we focus on time series segmentation techniques in the context of land cover change detection. The key contributions of this paper are:

(1) We propose a model-based time series segmentation algorithm inspired by a statistical event detection framework.
(2) We present a novel model-free change detection algorithm for detecting land cover change that is computationally simple, efficient, non-parametric and takes into account the inherent variability present in the remote sensing data. A key advantage of this technique is that it can be applied globally for a variety of vegetation without having to identify the right model for specific vegetation types.
(3) We evaluate the change detection efficacy of the proposed techniques on two data sets (i) simulated 16-day EVI time series containing phenological changes, and (ii) 16-day MODIS EVI time series for a region in North Carolina for which an independent land cover classification is available from National Oceanic and Atmospheric Administration (NOAA).
(4) We then illustrate the importance and the relative ability of different algorithms to account for the natural variation in the EVI data set due to different vegetation types, climate variability, geographic variability and errors in the data.

Organization of the Paper. Section 2 discusses previous work on segmentation techniques for land cover change. In Section 3, we present the two change detection algorithms. Section 4 describes the data used for experimentation. Section 5 presents the experimental evaluation with both simulated and real input data sets. Section 6 contains concluding remarks. Note that most figures in this paper are best seen in color.

2. Time Series Segmentation Techniques for Land Cover Change Detection: Related Work

In this section, we discuss the time series segmentation algorithms in the context of land cover change detection. In general, effective techniques for land cover change detection must be (i) scalable to handle large scale high resolution data sets; (ii) stable and robust to varying vegetation types; (iii) take into account noise and inherent variability present in the Earth Science data.

One commonly used segmentation-based approach divides a time series into multiple segments such that each segment can be approximately represented by a piecewise linear curve [13, 15]. The two key steps in this approach are to determine the best linear curve within a segment and to determine the number of segments in a time series. These techniques have been used in the remote sensing community for extracting phenology characteristics (e.g. timing of maximum of the growing season, length of growing season, onset of vegetation green-up) of the time series per year [7]. However, our work differs in that its objective is not to extract phenology characteristics but rather to identify the changes in the time series.

Below, we present the next two broad categories, model-based and model-free segmentation approaches and how they can be adapted for land cover change detection. We focus on identifying only one change in the time series though many of the techniques may be extended to find multiple changes. We specifically discuss the relative capabilities of these techniques to handle inherent variability and noise present in the data.

Model-based techniques involve fitting a model to a given time series. One such technique was proposed by Guralnik et al in [12]. It considers segmentation as a problem of either recognizing the change of parameters in the underlying model or the change of the most suitable model fit to the time series. It is an iterative algorithm that fits a model to a time segment, and uses a likelihood criterion to determine if the segment should be partitioned further. This approach is a top down strategy [15] which works by considering every possible partitioning of the time series and splitting it at the best location. Both the segments of the time series are then recursively partitioned in a similar way until a stopping criterion is reached. For single change point detection, the techniques aims at finding the first split. Therefore, in these techniques it is important to choose the correct model to represent the segments and an appropriate threshold as a stopping criterion. We adapt this technique for land cover change detection and evaluate it quantitatively. We also show that the choice of model plays a critical role in the performance of this algorithm.

Another *model-based* approach, Breaks for Additive Seasonal and Trend (BFAST) proposed recently by Verbesselt et al. [18] decomposes a time series into trend, seasonal and residual components. The time series is divided into segments such that intra-segment trend is constant and inter-segment trends are dissimilar. A trend breakpoint is associated with segment boundaries. The seasonal component is handled in a similar fashion. The focus of this work is on a paradigm for identifying multiple changes of different types, therefore we will not be comparing it directly in this paper. However, in the context of finding a single change in the seasonal pattern (which is the focus of our paper), BFAST is similar to the scheme presented in [12].

On the other hand, *model-free* time series segmentation algorithms do not assume any model for the time series but rather work directly with the data values. One such technique is the Recursive Merging algorithm proposed in [6] for land cover change detection. The algorithm starts with each year as the segmentation of original time series of length t, i.e with t/p (where p is the season length) segments. Next, the algorithm computes the cost of merging every adjacent pair of segments and iteratively merges the lowest cost pair. The process is repeated until two segments are left. The cost of merging can be computed in different ways, such as linear interpolation or linear regression. The algorithm also incorporates the notion of variability in the time series and is shown to be more effective than CUSUM [14] based change detection techniques and the change detection technique proposed by Lunetta et al. in [16]. In our paper, we will propose another novel model-free segmentation and compare it quantitatively with recursive merging and our adapted model-based algorithm.

3. Change Detection Techniques

This section describes the two change detection methods proposed in this paper. In Section 3.1 we propose a model-based segmentation technique inspired by a framework proposed in [12]. In Section 3.2 we propose a novel, simple and efficient model-free change detection algorithm which offers scalability and robustness to varying characteristics of time series across the globe.

All these approaches take as input the vegetation index time series and the annual season length for a location and give as output the corresponding change score and change point. The locations under study can be ranked according to the change score given by each algorithm. A good algorithm will give higher ranking to the locations that are more likely to have changed.

Following are the list of notations used in this paper. Let D be a data set with N land locations each of which has a time series of length T. The time series for a location corresponds to T 16-day EVI observations at that location. We also define the following notation:

p**:** season length (here it is 23)
Y**:** the number of years of data in the data set $=\frac{T}{p}$
n_i**:** an individual location $\quad n_{ij}$: EVI value at time j for the location n_i.
$b_{i1}, \ldots, b_{iy}$**:** list of annual cycles where, $b_{i1} = [n_{i,1},\, n_{i,2},\, \ldots,\, n_{i,23}]$, $b_{i2} = [n_{i,24},\, \ldots,\, n_{i,46}]$,

3.1. **Model-Based Segmentation Algorithm.** This approach follows a top-down segmentation strategy and is inspired from a framework proposed by Guralnik and Srivastava [12]. The technique follows an iterative algorithm that fits a model to a time segment, and uses a likelihood criterion to determine if the segment should be partitioned further, i.e. if it contains a new change-point. In other words, the likelihood criterion determines the statistical significance that a given time series should either be defined using a different set of model parameters or two different models. The need for two different models or a different set of parameters indicates that the time series contains a change point.

In Algorithm 1 we provide the general framework of the change detection scheme and provide specific details in the following paragraph.

Let p be the seasonal length
for each time series ts in a given dataset **do**
  Consider the entire time series $[n_{i,1},\, n_{i,2},\, n_{i,3},\, \ldots,\, n_{i,T}]$ as a single segment
  Choose a *model that best* fits the time series ts
  Calculate the *error of model* L from the original time series
  for each possible candidate timestamp $t = p \times j$, where $j \in [2, 3, \ldots, \frac{length(ts)}{p} - 2]$ **do**
    Divide the time series into two segments at t
    For each segment fit the *best model separately* and calculate the *individual errors* $-$ $L1$, $L2$
  end for
  Choose $\min(L1 + L2)$, which is the minimum of $L1 + L2$ over all possible values of t
  Score(S_i) of ts is $\frac{L - \min(L1+L2)}{L}$;
  Change Point for this time series ts is the index where $min(L1 + L2)$ occurs
end for

Algorithm 1: Model-based segmentation approach for time series.

Algorithm 1 has three key aspects which we address below:

(1) *Error computation between the model and the original time series of the segment:* In [12] the error for the model was calculated using residual sum of squares between the fitted model and the original time series. However, EVI time series contains noise due to cloud contamination which results in the sudden rise or fall of values in the time series. Since the residual sum of squares is sensitive to outliers, these spikes in the EVI data make the error computation less robust. Therefore, for EVI time series, we use the Manhattan distance between the model and the segment as the error value.
(2) *Choice of Model to fit the time series:* The choice of appropriate model plays a critical role in the performance of the scheme. There are two key properties that the model should possess in this framework (i) the model should follow the seasonality of the EVI vegetation time series data (ii) the model should not follow the change very well. For example, a piecewise polynomial model, which follows both the seasonality and change, cannot be used as it would

result in low error even if applied to a time series that is changed. Also, a non-seasonal model results in high values of L, $L1$ and $L2$ and thus a lower score even for a changed time series.

In this paper, we use a harmonic model which was inspired from the work by Verbesselt et al. [18] for estimating an EVI time series. The harmonic model follows the seasonality well and is less sensitive to short term data variations and noise. The value of parameter K used in the analysis is 3. We refer to this scheme as *HM-Variability* in this paper, where HM stands for Harmonic Model.

(3) *Score for the Time Series:* In the original scheme, at every iteration, the value of the likelihood criterion was calculated until it fell below a certain threshold. However, since the focus is on a single change, we use the maximum value of the likelihood criterion obtained in the first iteration as the change score. Normalization of the likelihood estimation by L in the above scheme models the inherent variability of the time series. To evaluate the ability of *HM-Variability* to model variability, we evaluate a variation of it that does not perform the normalization step. We refer to that scheme as *HM-NoVariability*.

3.2. **Model-Free Segmentation Algorithm:** The two key characteristics of this algorithm are (i) the technique does not assume any model for the time series but rather works directly with the data values. It can therefore be applied to any periodic data without having to choose an appropriate model (ii) the technique introduces a new method to incorporate the notion of variability in the time series due to both noise in the data and climate variations.

The algorithm assumes that each time series undergoes a maximum of one phenological pattern change. In particular, it assumes that a changed time series follows a certain pattern for the first few years and then follows a different pattern for the next few years. For a non-changed pixel, its time series follows the same pattern throughout its time period. There is a notion of pattern for each annual segment. This technique does not use any model and is non-parametric.

The key idea of the proposed algorithm is to find two continuous segments in the time series such that the annual years (objects) within each segment are very similar to each other while being significantly different from the objects across the segments. The boundary of the segments represents the change point in the time series. To model the similarity and differences between the objects for each segment we calculate two terms: *Cohesion* and *Separation*. The cohesion of a segment is defined as an average of the pairwise distance of all annual years within the segment. Cohesion for the time series of a pixel is defined as an average of the cohesion of both the segments (see Figure 2). The value of cohesion gives an estimate of the natural variability within the time series. Higher values of cohesion indicate higher natural variability of the time series since it means that the distances between the years in the same segment are also high. For example, the value of cohesion for a time series with no noise or fluctuations would be zero since the annual cycles would look exactly same within each segment. Likewise, the separation between two segments can be measured by the sum of the distances from objects in one segment to objects in the other segment. The value of separation indicates how distinct or well-separated the segments are to each other. The combination of cohesion and separation values indicates the amount of change in the time series with respect to the natural variation. In Algorithm 2, we describe in detail how every pixel is assigned a score and a change point.

$$C(i,t) = \frac{\frac{\sum_{p=1:t}\sum_{q=1:t} M(p,q)}{t^2-t} + \frac{\sum_{p=t+1:Y}\sum_{q=t+1:Y} M(p,q)}{(Y-t)^2-(Y-t)}}{2} \tag{1}$$

$$S(i,t) = \frac{\sum_{p=1:t}\sum_{q=t+1:Y} M(p,q)}{(Y-t)\times t} \tag{2}$$

Note that we assume that the change points occur no earlier than the end of second year and no later than the second to last year since we want at least two annual years to be present in each segment to account for inter-annual variability.

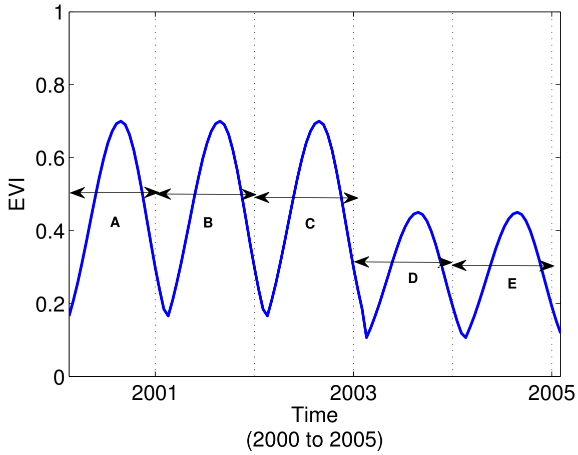

(a)

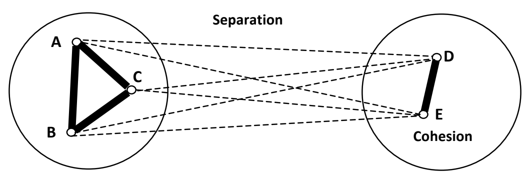

(b)

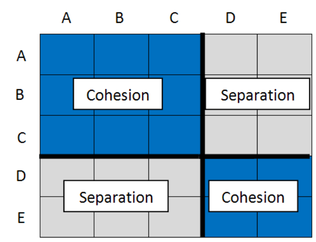

(c)

FIGURE 2. Illustration of Cohesion and Separation. (a) Time series with different years A,...,E; (b) Two different circles containing 3 and 2 points (shown as small circles) represent two segments. The dark edges represent the cohesion and the dotted lines represent separation between the segments; (c) Dissimilarity matrix constructed by using the pairwise distances between years.

for each time series ts in a given dataset **do**
  Create a dissimilarity matrix M for the time series
  Each entry $M(q, r)$ in the matrix contains the distance between the annual segments b_{iq}, b_{ir}
  We use Manhattan distance between the vectors b_{iq} and b_{ir}
  for each possible candidate timestamp $t = p \times i$, where $i \in [2 \cdots \frac{length(ts)}{p} - 2]$ **do**
    Cohesion $(C(i, t))$ with respect to t is calculated as in Equation 1
    Separation $(S(i, t))$ with respect to t is calculated as in Equation 2
    $Score(i, t) = S(i, t) - C(i, t)$
  end for
  $ChangeScore(i) \equiv max_t Score(i, t)$
  Change Point of this time series ts is the index where $max_t Score(i, t)$ occurs
end for

Algorithm 2: Our proposed model-free segmentation algorithm

The key aspect of this algorithm is the use of values of cohesion and separation to distinguish a real change from the natural variability of the time series. Using Figure 3, we illustrate how the distance matrix looks for different types of series and the capability of the method to use all of the existing information to incorporate variability and assign change scores. Consider the following:

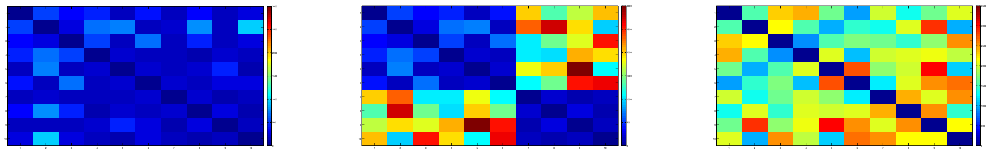

(a) Unchanged time series with low variability

(b) Changed time series

(c) Unchanged time series with high variability

FIGURE 3. Dissimilarity matrices for different kinds of time series. The blue values represent low values while the red and yellow values are higher

(1) A time series with no change and low variability: The dissimilarity matrix for such a time series is shown in Figure 3(a). Since each annual segment for such a time series would be very similar, the dissimilarity matrix consists of low values which results in low cohesion and separation, resulting in a overall lower score.
(2) A stable time series with a change: The typical dissimilarity matrix for such a time series is shown in Figure 3(b) Notice that the separation values are high and the cohesion values are low which results in a high score. Visually, notice that dissimilarity matrix has a roughly block diagonal structure since the time series have well-seperated segments.
(3) A highly variable time series with no change: The dissimilarity matrix shown in Figure 3(c) consists of all high values. If we consider only the separation between any two segments, we would obtain a high score for the time series and wrongly label it as change. However, if we consider the cohesion between the time series and the relative difference between cohesion and separation, the time series would be given a low score since all values are relatively similar. Visually also, there is no block-diagonal structure observed in the dissimilarity matrix signifying that the time series does not have well-seperated segments.

The above discussion illustrates the importance of including measures of variability in the analysis of vegetation index data set to effectively distinguish between an unusual event and an event within the normal range of variability. In this paper, we refer to the scheme using only separation as *MF-NoVariability* and using the difference of separation and cohesion i.e., $S(i,t) - C(i,t)$ as *MF-Variability*.

Another way to handle variability in the time series is to examine the distribution of the pairwise distance values between the objects in the same segment and object across the segment. In this paper we use the t-statistic as the scoring function. We refer to this scheme as *MF-T-stat*. The scoring function in Algorithm 2 is replaced by the score below:

$$Score = \frac{tstatistic(S(i,t), C(i,t)_{Seg1}) + tstatistic(S(i,t), C(i,t)_{Seg2})}{2}$$

4. DATA AND EVALUATION METHODOLOGY

Below, we provide details of the simulated and the MODIS EVI data sets used for evaluation.

4.1. **Simulated 16-day EVI Time Series:** Simulated EVI time series are generated by summing simulated seasonal and noise components. This procedure was adapted from [18]. The seasonal component is created using an asymmetric Gaussian function (Equation 3) for each season. Two different kinds of seasonal cycles are created by using $x \in [1, p]$ for single cycles per year and $x \in [1, \frac{p}{2}]$ for double cycles per year, where p is the season length.

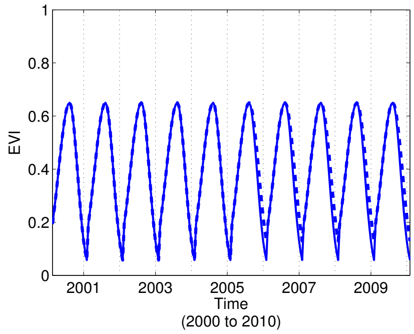

FIGURE 4. Seasonal change by changing the C_1 from 5 0(–) to 75 (dashed): c_2=100, b=12

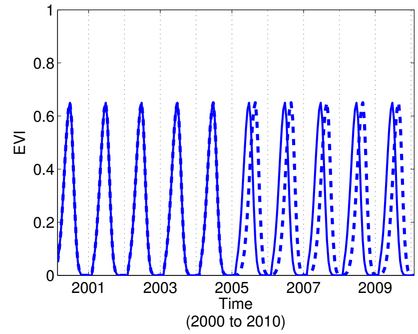

FIGURE 5. Seasonal change by changing b from 9 (–) to 13 (dashed): c_1 = 10, c_2 = 25

$$f(x) = \begin{cases} ae^{\frac{-(b-x)^2}{c_1}} & x \geq b, \\ ae^{\frac{-(b-x)^2}{c_2}} & \text{otherwise.} \end{cases} \tag{3}$$

The parameters a and b determine the amplitude and the position of maximum or minimum with respect to the independent time variable t, while c_1 and c_2 determine the width of the left and the right hand side, respectively.

In addition to the seasonal component, the following two noise components were generated (i) *Noise_Seasonal*: Simulates the inter annual seasonal variability observed in the EVI values due to climate variations and was generated using a random number generator that follows a uniform distribution over a pre-defined range (ii) *Noise_Spike*: A noise component that was added to a pre-defined number of time stamps in each time series to simulate cloud contamination. The value of the noise component also followed a uniform distribution between a pre-defined range.

The time series were generated in the following manner:

(1) *Unchanged Time Series:* The same values of the parameters are used in Equation 3 for all individual years. Both noise components were added using the method as described above.
(2) *Changed Time Series:* The changed time series are constructed by changing the parameters within a single time series after a certain year which is chosen randomly between years 2 and 8. Different parameters impact the time series in a different way. For example, Figure 4 illustrates a pattern change introduced from fifth year onwards by changing c_1 from 10 to 100 while keeping all the other parameters fixed. The change in different parameters corresponds to different land cover changes. For example, a change in only the amplitude might represent a degradation of crop productivity or a change in only c_1 or c_2 might indicate a different cropping pattern. However, to easily compare the relative performance of different algorithms we only change the parameter b in the two different segments. Figure 5 shows the effect of changing the parameter b in the two segments.

Using the above procedure different data sets were created which differed in the amount of noise:

DS1: It is a combination of multiple data sets (Table 1) which have similar amplitude range but vary in the levels of noise. This was to simulate the areas with similar vegetation patterns but different characteristics of noise due to geographic locations, climate patterns etc. All the data sets used have the same number of changed and non-changed time series and contained both single & double cycled time series in equal proportion.

DS2: This data set was created to simulate changes occuring in different vegetation phenologies having different noise levels and extent of changes. The constructed data set is the combination of

Name	Amplitude	$Noise_{Seasonal}$	$Noise_{Spike}$	% of time stamps	Changed	Non-Changed
DS-N1	[3000,7000]	[-500,500]	[1200,1500]	10	2000	20,000
DS-N2	[3000,7000]	[-500,500]	[1700,2000]	30	2000	20,000
DS-N3	[3000,7000]	[-1000,1000]	[1200,1500]	10	2000	20,000
DS-N4	[3000,7000]	[-1000,1000]	[1700,2000]	30	2000	20,000
DS-N5	[3000,7000]	[-1500,1500]	[1200,1500]	10	2000	20,000
DS-N6	[3000,7000]	[-1500,1500]	[1700,2000]	30	2000	20,000

TABLE 1. Summary of different data sets used to create data set **DS1**

the data sets shown in Table 2. Notice that it contains two data sets: DS-N7 with higher amplitude & higher levels of noise and DS-N8 with lower amplitude & lower levels of noise.

Name	Amplitude	$Noise_{Seasonal}$	$Noise_{Spike}$	% of time stamps	Changed	Non-Changed
DS-N7	[3000,7000]	[-1500,1500]	[1700,2000]	30	2000	20,000
DS-N8	[1000,1500]	[-500,500]	[1200,1500]	10	2000	20,000

TABLE 2. Summary of different data sets used to create data set **DS2**

DS3: These data sets were created to illustrate the importance of choosing an appropriate model in the model based change detection algorithm. DS3 consists of data sets shown in Table 3. DS-N9 is constructed using Assymetric Gaussian function as in Equation 3. DS-N10 is however constructed using Wigner semicircle distribution model as in Equation 4. Changed time series are constructed by changing the values of $R \in [6, 11]$. Both these data sets have 2,000 changed and 20,000 nonchanged time series.

$$f(x) = \begin{cases} \frac{2}{\pi R^2}\sqrt{(R^2 - x^2)} & -R < x < R, \\ 0 & \text{otherwise.} \end{cases} \tag{4}$$

Name	Model	Amplitude	$Noise_{Seasonal}$	$Noise_{Spike}$
DS-N9	Asymmetric Gaussian	[8000,12000]	[-500,500]	[1200,1500]
DS-N10	Wigner semicircle	[8000,10000]	[-1500,1500]	[2000,2500]

TABLE 3. Summary of data set **DS3**

4.2. **16-day MODIS EVI Time Series:** The specific vegetation-related variable used in this analysis was the Enhanced Vegetation Index (EVI) product that serves as a surrogate for the amount of vegetation for a pixel; and is measured by the moderate resolution imaging spectroradiometer (MODIS) instrument. In this paper, the temporal coverage of the data is from the time period February 2000 – February 2010.

We selected a region in North Carolina containing 48,025 pixels of 250m resolution between North 35.99–35.3 and West 76.5–77. We refer to this data set as DSNC. This region was chosen because it is known to have a variety of changes in land cover over the past 10 years. Also, a reasonably good quality land cover classification map of this region is available from NOAA [4] at 30m resolution for 2001 and 2006 that can be used for validation. Each 250m pixel was assigned a set of 30m pixels based on the nearest neighbor and a 250m pixel was considered a change if a certain threshold (10%

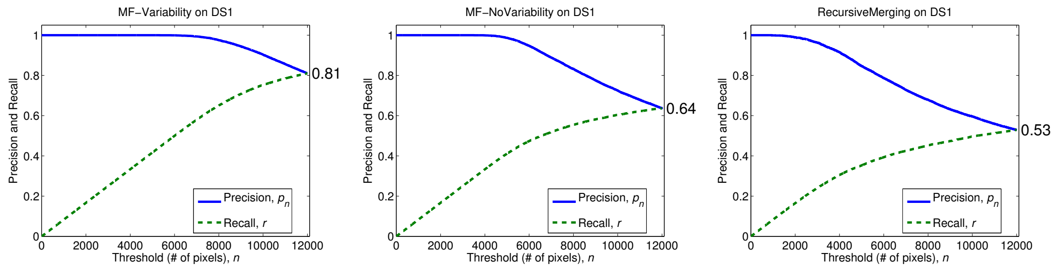

(a) (b) (c)

FIGURE 6. Precision-Recall curve for DS1, Blue curve is Precision, Green curve is Recall. x-axis represents the number of pixels (a) *MF-Variability*; (b) *MF-NoVariability*; (c) *RecursiveMerging*

in our analysis) of the 30m pixels within that 250m pixel had different land cover labels in 2001 and 2006. Using this threshold 7,367 pixels were considered changed. More details of the ground truth generation are provided in the technical report [9].

4.3. **Evaluation Methodology.** Assume that for a time series data set D with N pixels, the change detection technique returns a list of *change scores* of length N, where each change score is a measure of the degree of change for the corresponding pixel. We also have a validation data set which consists of true labels for each of the pixels; let M be the *total* number of actual changes as determined by the validation data set. To evaluate the performance of a given change detection algorithm at rank n, we count the number of true changes in the top n ranked pixels of the sorted change scores of all the pixels, where n is the number of actual changes ($1 \leq n \leq M$). Let TP_n be the number of actual disturbances in the top n predicted disturbances, and FP_n be the number of pixels that are in the top n portion but are not actual disturbances.

We evaluate performance by examining the *sorted* list of change scores. The performance metrics are defined as follows:

$$\text{Precision, } p_n = \frac{TP_n}{TP_n + FP_n} \qquad\qquad \text{Recall, } r_n = \frac{TP_n}{M}$$

Note that as n increases, p_n will tend to decrease and r_n will increase. One specific value of interest is the one when n is equal to the number of changed pixels (validation data). At this value of n, $p_n = r_n$. Also, if the change detection algorithm does a perfect job of identifying changes, then p_n will remain at 1 upto this value of n and then start to drop for increasing values of n and r_n will linearly increase from 0 to 1 and then stay at 1 for larger values of n.

5. Experimental Results

5.1. **Observations on Simulated Data Sets:** Below we present precision and recall curves for different algorithms on DS1, DS2, DS3 and DSNC. We particularly focus on the relative capabilities of the algorithm to model natural variation. We also present results to illustrate the dependence of model based algorithms on the choice of model.

5.1.1. ***MF-Variability Significantly Outperforms MF-NoVariability and RecursiveMerging for DS1:***

The precision and Recall curves in Figure 6 for DS1 shows that *MF-Variability* significantly outperform *MF-NoVariability* and *RecursiveMerging*. The primary reason is that since dataset DS1 consists of time series with varying levels of variability (noise), the change detection algorithm must take into account the change with respect to the natural variation. Since *MF-NoVariability* does not depend on the value of cohesion, it is not able to model the natural variation in the time series.

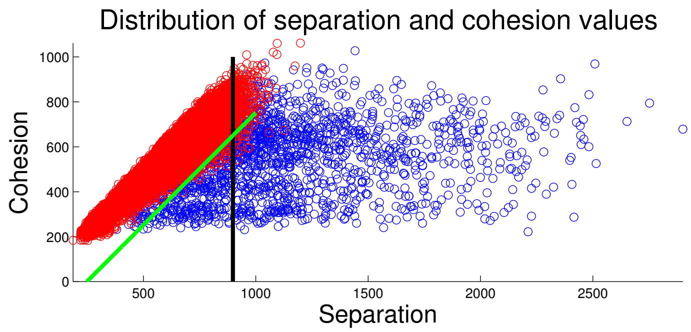

FIGURE 7. Cohesion-Separation value distribution for changed and unchanged pixels. The changed time series are represented by blue circles while the unchanged pixels are represented by red circles.

To illustrate the advantage of subtracting the cohesion from separation in *MF-Variability*, in Figure 7 we show the scatter plots of the Separation and Cohesion values for a random sample of 2,000 changed (blue circles) and 20,000 non-changed time series (red circles) from DS1. The vertical line in black shows the constant *MF-NoVariability* score of 898 and the oblique line in green shows the constant score of *MF-Variability* score of 248. Points lying to the right half of these lines will have scores higher than the respective line. These scores are chosen because they give similar number of changed events. From the Figure 7, we notice that *MF-NoVariability* will make more errors as compared to *MF-Variability* by incorrectly labelling a few unchanged time series as changed.

The discussion above illustrates that the notion of variability is important to incorporate in the change detection algorithm. Using the value of cohesion as an indicator of the natural variation, *MF-Variability* is able to significantly improve the results.

5.1.2. ***MF-T-stat outperforms MF-Variability for DS2:***

Figure 8(a) shows the precision and recall curve for *MF-Variability* on DS2. Recall that DS2 consists of two different kinds of vegetation patterns: (i) time series with higher amplitude and higher levels of noise (DS-N7) (ii) time series with lower amplitude and lower level of noise (DS-N8). Table 4 shows the number of true and false positives from the individual data sets DS-N7 and DS-N8 when *MF-Variability* is used on DS2. It is seen that only 325 points out of 2,000 changed points are recalled from the data set DS-N8. Also, notice that almost all the false positives are from the dataset DS-N7. This illustrates that because of the higher levels of noise present in DS-N7 and smaller number of changes in DS-N8, *MF-Variability* gives a higher score to unchanged time series in DS-N7 than compared to changed time series in DS-N8. It is therefore important to design a scoring mechanism which takes into account the difference in variation observed in the time series due to different phenological characteristics. As discussed in Section 4, Figure 8(a) illustrates how *MF-T-stat* models the variance of the distribution in cohesion and separation values to significantly improve the results. Table 5 further illustrates that the *MF-T-stat* is able to recall many more points from DS-N8 as compared to *MF-Variability*.

5.1.3. ***HM-Variability outperforms HM-NoVariability for DS2:***

Figure 9 shows the precision recall curve for *HM-Variability* and *HM-NoVariability*. It is seen that *HM-Variability* significantly outperforms *HM-NoVariability*. The primary reason for better performance of *HM-Variability* is similar to as explained above for the comparison of *MF-T-stat* and *MF-Variability*. The normalization step in *HM-Variability* helps to model the difference in variability of the two different combined data sets.

5.1.4. ***Model Choice plays a critical role in the performance of Model Based Algorithm:***

To illustrate the importance of model choice, we show results on DS3 which consists of time series generated from two different models: asymmetric Gaussian (DS-N9) and Wigner Semicircle

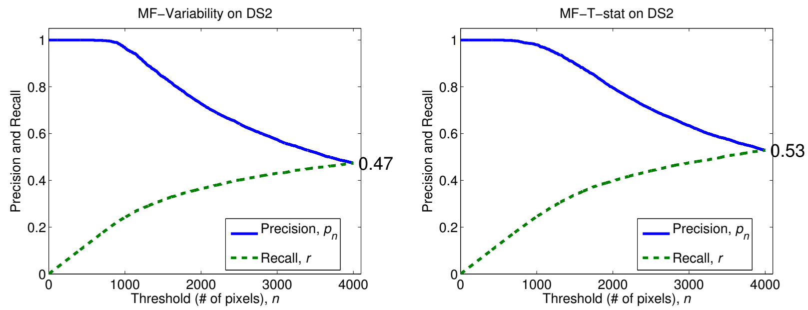

(a) (b)

FIGURE 8. Precision-Recall curve for DS2 (a) *MF-Variability* (b) *MF-T-stat*

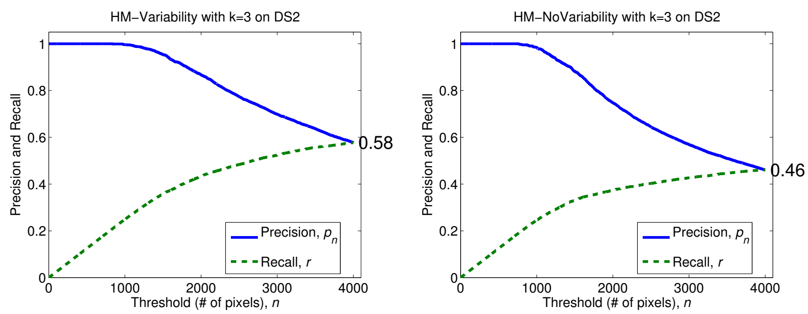

(a) (b)

FIGURE 9. Precision-Recall curve for DS2 (a) *HM-Variability* (b) *HM-NoVariability*

TP or FP	**DS-N7**	**DS-N8**
TP	1569	325
FP	2102	4

TABLE 4. *MF-Variability*

TP or FP	**DS-N7**	**DS-N8**
TP	1328	828
FP	925	918

TABLE 5. *MF-T-stat*

Distribution (DS-N10), as mentioned in Section 4. On this data set, *MF-Variability* significantly outperforms *HM-Variability* as shown in Figure 10. The primary reason is that since the harmonic model used in *HM-Variability* does not appropriately model the time series in DS-N10, the error computation is not accurate. In particular, the error between the fit and the original time series is particularly high for time series in DS-N10, resulting in lower score being assigned to such time series due to the normalization step in *HM-Variability.* This is also represented in the number of true and false positives detected by the algorithms for the individual data sets present in DS3 (shown in Table 6 and Table 7). Note that only a few (767 out of 2,000) changed points are recalled from the data set DS-N10. Also, notice that all of the false positives are from the data set DS-N9. Therefore, the choice of the model in model-based algorithm is critical to its performance. On the other hand, *MF-Variability* does not require any knowledge of model or choice of parameter and therefore is robust to different phenologies and characteristics of time series globally. This is one of the key properties of the *MF-Variability* algorithm for its application in global land cover change detection.

TP or FP	DS-N9	DS-N10
TP	2000	767
FP	1233	0

Table 6. *HM-Variability*

TP or FP	DS-N9	DS-N10
TP	2000	1637
FP	0	363

Table 7. *MF-Variability*

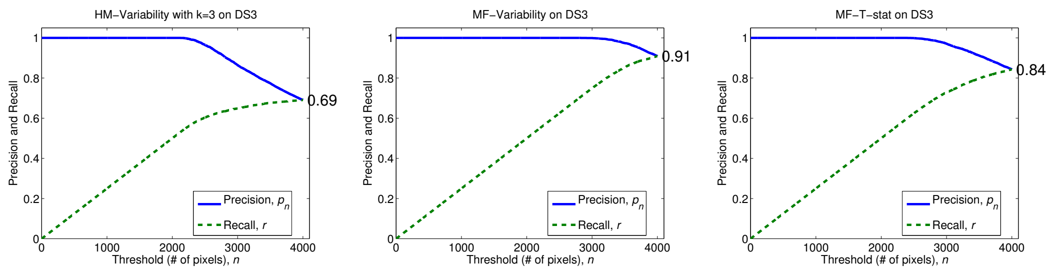

Figure 10. Precision-Recall curves for DS3 (a) *HM-Variability*; (b) *MF-Variability*; (c) *MF-T-stat*

5.2. **Observations on real dataset: DSNC:.** Figure 11 shows the precision recall curve for different algorithms on DSNC. It is observed that none of the algorithms perform very well on this dataset. This is primarily due to various issues associated with the validation data set which complicates the evaluation. First, the resolution difference between the label dataset (30m) and the MODIS EVI data set (250m) results in inaccuracy in assigning the proper set of labels to each 250m pixel. Also, determining the threshold for the number of 30m pixels required to have changed for each 250m pixel to be considered as change is challenging. For example, though a conversion of 10% of 30m pixels within a 250m pixel from forest to barren land could be strongly reflected in the 250m EVI signal, a 10% conversion from forest to pasture might not be reflected. Additional challenges arise from the inability of the EVI signal to distinguish between some particular land cover types. A pixel classified as Secondary Forest in 2001 and Mixed Forest in 2006 is considered changed according to the validation data set but might not show a perceptible change in its EVI signal and thus would not be detected by the change detection algorithm. Conversely, certain changes such as double cropping to single cropping cycles which are clearly reflected in the EVI signal are not considered change according to the ground truth because they have the same LCC label. Such pixels detected by the algorithm are considered as false positives by the evaluation methodology and thus reduces the observed performance of the algorithms.

Despite these challenges, note that all the algorithms still do significantly better than the random curve shown in Figure 11. Also, it is observed that *MF-Variability* performs the best and significantly better than *MF-NoVariability*. However, it is difficult to make precise statements about the relative performance given the uncertainity associated with the validation labels.

6. Conclusion and Future Work

In this paper, we presented two time series segmentation techniques that can be used to identify the pattern changes in the vegetation index time series. The results of this study also demonstrate the importance of modeling the natural variation in the time series for accurately estimating the significance of the change in the EVI signal. Both the techniques significantly outperformed another recently proposed technique by Boriah et al [6]. The proposed model-based segmentation algorithm was shown to be sensitive to the choice of model, however the model-free segmentation algorithm requires no model and gives comparable or better results. The proposed model-free segmentation

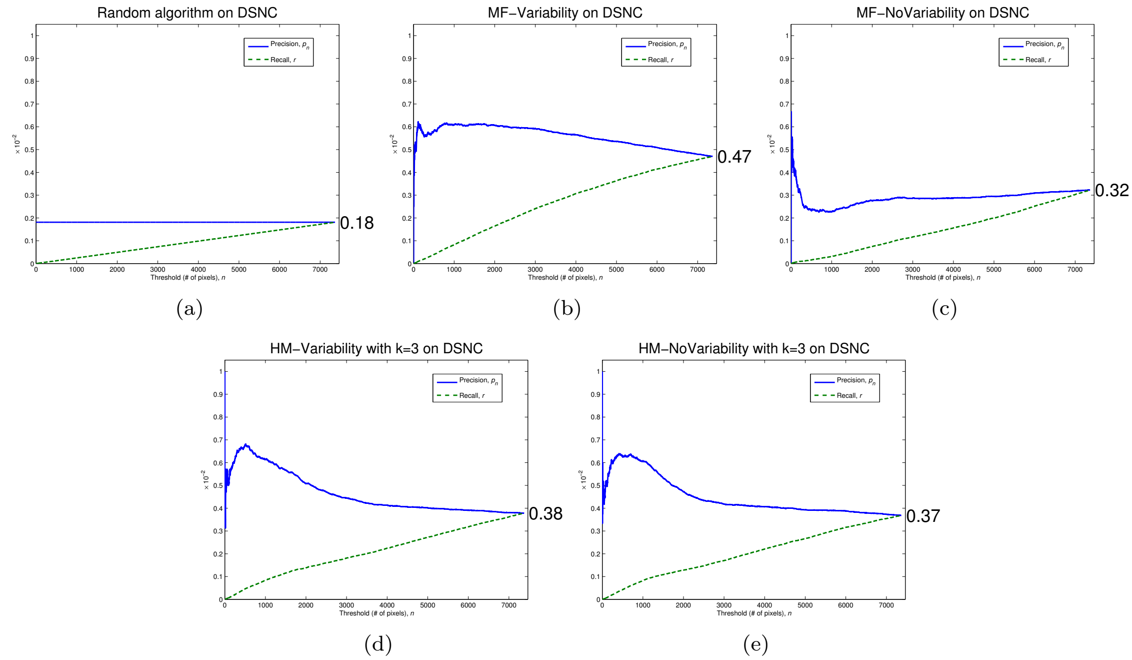

FIGURE 11. $x - axis$ shows the number of pixels considered; $y - axis$ shows the precision (range 0-1) and recall (range 0-1); Precision-Recall curves on real data set (a) *Random Algorithm* (b) *MF-Variability*; (c) *MF-NoVariability*; (d) *HM-Variability*; (e) *HM-NoVariability*

algorithm has been applied globally at 1km EVI to detect various land cover changes [10] such as forest to farmland conversions, change in cropping patterns, urbanization and the results are publicly available via the online platform ALERTS [1]. The results indicate the ability of the algorithm to provide rapid, inexpensive, robust, scalable and precise detection of land use change [2].

The proposed algorithms assume that only one pattern change occurs in the time series. However, the ability to find multiple changes becomes critical as the length of the time series increases with the continuous collection of satellite data. Therefore the existing techniques ought to be extended to detect multiple changes. This could be challenging since the presence of multiple change points might hinder the effective detection of the first change point using the top down segmentation approaches. In addition, the techniques need to be adapted to discover changes even in the presence of other changes such as gradual or abrupt drops. BFAST, a recently proposed technique [18] outlines a framework to detect such changes, however the technique is computationally expensive and hence not scalable for global application. The proposed techniques could be extended using similar frameworks to detect such changes. Also, our techniques assume that the pattern changes occur at the yearly boundaries which is not always true in the land cover change domain.

Acknowledgements

The research described here is supported by the Planetary Skin Institute, NSF Grant IIS-0713227, NSF Grant IIS-0905581, NASA grant NNX09AL60G and the University of Minnesota MN Futures Program. Access to computing facilities was provided by the University of Minnesota Supercomputing Institute. We are grateful to Ivan Brugere, Xi Chen, Divya Alla, Sruthi Reddy Vangala and Yashu Chamber for their helpful comments.

References

[1] Planetary Skin.
`http://www.ourplanetaryskin.org/`.

[2] Monitoring Forests: Seeing the world for the trees.
`http://www.economist.com/node/17730208/`.

[3] Land Processes Distributed Active Archive Center.
`http://edcdaac.usgs.gov`.

[4] Coastal Change Analysis Program Regional Land Cover.
`http://www.csc.noaa.gov/digitalcoast/data/ccapregional/`.

[5] S. Boriah. *Time Series Change Detection: Algorithms for Land Cover Change*. PhD thesis, University of Minnesota, 2010.

[6] S. Boriah, V. Kumar, M. Steinbach, C. Potter, and S. Klooster. Land cover change detection: A case study. In *KDD 2008: Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 857–865, 2008.

[7] V. Chandola, D. Hui, L. Gu, B. Bhaduri, and R. R. Vatsavai. Using time series segmentation for deriving vegetation phenology indices from modis ndvi data. In *Proceedings of the 2010 IEEE International Conference on Data Mining Workshops*, ICDMW '10, pages 202–208, Washington, DC, USA, 2010. IEEE Computer Society.

[8] P. Coppin, I. Jonckheere, K. Nackaerts, B. Muys, and E. Lambin. Digital change detection methods in ecosystem monitoring: a review. *International Journal of Remote Sensing*, 25(9):1565–1596, 2004.

[9] A. Garg, L. Manikonda, S. Kumar, V. Krishna, S. Boriah, M. Steinbach, V. Kumar, D. Toshniwal, C. Potter, and S. Klooster. Pre-processing of the validation data used in the paper titled "Model-Free Time Series Segmentation Approach for Land Cover Change Detection". *Technical Report 11-017, Computer Science Department, University of Minnesota*, 2011.

[10] A. Garg, V. Mithal, Y. Chamber, I. Brugere, V. Chaudhari, M. Dunham, V. Krishna, S. Krishnamurthy, S. Vangala, S. Boriah, M. Steinbach, V. Kumar, A. Cho, J. Stanley, T. Abraham, J. C. Castilla-Rubio, C. Potter, and S. Klooster. GOPHER: Global observation of planetary health and ecosystem resources. In *IGARSS 2011: Proceedings of the IEEE Geoscience and Remote Sensing Symposium*, 2011.

[11] L. Giglio, G. R. van der Werf, J. T. Randerson, G. J. Collatz, and P. Kasibhatla. Global estimation of burned area using modis active fire observations. *Atmospheric Chemistry and Physics*, 6(4):957–974, 2006.

[12] V. Guralnik and J. Srivastava. Event detection from time series data. In *Proceedings of the fifth ACM SIGKDD international conference on Knowledge discovery and data mining*, KDD '99, pages 33–42, New York, NY, USA, 1999. ACM.

[13] D. M. Hawkins. Point estimation of the parameters of piecewise regression models. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 25(1):pp. 51–57, 1976.

[14] K. Jan, B. F. Paulo, and S. Peter. Cumulative sum charts - a novel technique for processing daily time series of modis data for burnt area mapping in portugal. In *In MultiTemp 2007: International Workshop on the Analysis of Multi-temporal Remote Sensing Images*, pages 1–6, 2007.

[15] E. Keogh, S. Chu, D. Hart, and M. Pazzani. Segmenting time series: A survey and novel approach. In *In an Edited Volume, Data mining in Time Series Databases. Published by World Scientific*, pages 1–22. Publishing Company, 1993.

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

[17] V. Mithal, A. Garg, S. Boriah, M. Steinbach, V. Kumar, C. Potter, S. Klooster, and J. C. Castilla-Rubio. Monitoring global forest cover using data mining. *ACM Trans. Intell. Syst. Technol.*, 2:36:1–36:24, July 2011.

[18] J. Verbesselt, R. Hyndman, A. Zeileis, and D. Culvenor. Phenological change detection while accounting for abrupt and gradual trends in satellite image time series. *Remote Sensing of Environment*, 114(12): 2970 – 2980, 2010.