A COMPARATIVE STUDY OF ALGORITHMS FOR LAND COVER CHANGE

SHYAM BORIAH*, VARUN MITHAL*, ASHISH GARG*, VIPIN KUMAR*, MICHAEL STEINBACH*, CHRIS POTTER**, AND STEVE KLOOSTER***

Abstract. Ecosystem-related observations from remote sensors on satellites offer huge potential for understanding the location and extent of global land cover change. This paper presents a comparative study of three time series based algorithms for detecting changes in land cover. The techniques are evaluated quantitatively using forest fire ground truth from the state of California for 2000–2009. On relatively high quality data sets, all three schemes perform reasonably well, but their ability to handle noise and natural variability in the vegetation data differs dramatically. In particular, one of the algorithms significantly outperforms the other two since it accounts for variability in the time series.

1. Introduction

The climate and earth sciences have recently undergone a rapid transformation from a data-poor to a data-rich environment. In particular, climate and ecosystem related observations from remote sensors on satellites, as well as outputs of climate or earth system models from large-scale computational platforms, provide terabytes of temporal, spatial and spatio-temporal data. These massive and information-rich datasets offer huge potential for advancing the science of land cover change, climate change and anthropogenic impacts.

One important area where remote sensing data can play a key role is in the study of land cover change. Specifically, the conversion of natural land cover into human-dominated cover types continues to be a change of global proportions with many unknown environmental consequences. In addition, being able to assess the carbon risk of changes in forest cover is of critical importance for both economic and scientific reasons. In fact, changes in forests account for as much as 20% of the greenhouse gas emissions in the atmosphere, an amount second only to fossil fuel emissions.

Thus, there is a need in the earth science domain to systematically 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. Land cover conversions include tree harvests in forested regions, urbanization, and agricultural intensification in former woodland and natural grassland areas. These types of conversions also have significant public policy implications due to issues such as water supply management and atmospheric CO_2 output. In spite of the importance of this problem and the considerable advances made over the last few years in high-resolution satellite data, data mining, and online mapping tools and services, end users still lack practical tools to help them manage and transform this data into actionable knowledge of changes in forest ecosystems that can be used for decision making and policy planning purposes.

For ecosystem data, change detection is the process of identifying changes in the cover type and/or human use of the Earth. Examples include the conversion of forested land to barren land (possibly due to deforestation or a fire), grasslands to golf courses and farmland to housing developments. There is a large body of research in change detection using remotely sensed image data. Most previous change detection studies primarily rely on examining differences between two or more satellite images acquired on different dates [9]. However, these techniques have well-known limitations (as

* University of Minnesota, <sboriah,mithal,ashish,kumar,steinbac>@cs.umn.edu
** NASA Ames Research Center, chris.potter@nasa.gov
*** California State University Monterey Bay, skglooster@gaia.arc.nasa.gov.

discussed in Section 2) and are suitable for use in relatively small areas or to describe changes in specific categories of interest [8, 13, 20, 21] because they are inherently unsuited for global analysis.

More recently, several time series change detection techniques have been explored in the context of land cover change detection. Lunetta et al. [17] presented a change detection study that uses MODIS data and evaluated its performance for identifying land cover change in North Carolina. Kucera et al. [15] describe the use of CUSUM for land cover change detection. However, no qualitative or quantitative evaluation was performed. The Recursive Merging algorithm proposed by Boriah et al. [5] follows a segmentation approach to the time series change detection problem and takes the characteristics of ecosystem data into account. They provide a qualitative evaluation using MODIS EVI (Enhanced Vegetation Index) data for the state of California and MODIS FPAR (Fraction of Photosynthetically Active Radiation) data globally.

In this paper, we investigate the performance of these three techniques and their variations for the task of land cover change detection. In particular, we present a quantitative assessment of these techniques using the forest fire ground truth data in California and analyze the key characteristics of each technique that impact their suitability for land cover change detection problem.

1.1. **Key Contributions.** The key contributions of this paper are as follows:

- We systematically study the three algorithms (and their variations) for land cover change detection. We quantitatively evaluate their performance using forest fire ground truth from 2000—2009 for the state of California.
- We compare the three algorithms and their variations in their ability to handle variability inherently present in Earth Science data.

1.2. **Organization of the Paper.** We motivate the land cover change detection problem and discuss previous work in Section 2. In Section 3, we present the three change detection algorithms studied in this paper. Section 4 presents the experimental evaluation with multiple input data sets, and provides a discussion of the results. Section 5 contains concluding remarks. Note that most figures in this paper are best seen in color.

2. Time Series-based Land Cover Change Detection: Background and Related Work

There is an extensive literature on time series change detection that can, in principle, be applied to the land cover change detection problem. Time series based change detection has significant advantages over the comparison of snapshot images of selected dates since it can take into account information about the temporal dynamics of landscape changes. In these schemes, detection of changes is based on the pattern of spectral response of the landscape over time rather than the differences between two or more images collected on different dates. Therefore, additional parameters such as the rate of the change (e.g. a sudden forest fire vs. gradual logging), the extent, and pattern of regrowth can be derived. By contrast, for image-based approaches, changes that occur outside the image acquisition windows are not detected, it is difficult to identify when the changes occurred, and information about ongoing landscape processes cannot be derived. For illustration, Figure 1 shows an example of a land cover change pattern that is typically of interest to Earth Scientists. The time series shows an abrupt jump in EVI in 2003. The location of the point corresponds to a new golf course, which was in fact opened in 2003. Changes of this nature can be detected only with high resolution data.

Time series change detection, in general, is an area that has been extensively studied in the fields of statistics [12], signal processing [11] and control theory [16]. However, many of these techniques are not suitable for the land cover change detection problem primarily because they are not scalable or they are unable to take advantage of the inherent structure present in earth science data. For example, the major mode of behavior in the vegetation signal is seasonality, i.e., the natural seasonal growing cycle is a dominant characteristic of a time series and this intrinsic seasonality should not

Source: Google Maps.

Figure 1. This figure shows an example of a change point in the San Francisco Bay Area which corresponds to a new golf course constructed in Oakland, CA. This golf course was built in 2003, which corresponds to the time step at which the time series exhibits a change.

itself be called a change. In addition, there exists an inherent natural variability and noise in the earth science data because of local weather, geography, and atmospheric conditions. Additional challenges in global land cover change studies include the massive data size, high degree of geographic/inter-region variation, missing data, disparate land cover types, and the large variety of changes that can occur. There are three key approaches to time series change detection:

- **Parameter change:** In this setting, the time series is expected to follow a particular distribution and any significant departure from the distribution is flagged as a change. Fang et al. [10] presented a parameter change based approach for land cover change detection. CUSUM (and its variants) is the most well-known technique of the parameter change approach.
- **Segmentation:** The goal of the segmentation problem is to partition the input time series into homogeneous segments (the subsequence within a segment is contiguous). Segmentation is essentially a special case of change detection since by definition, successive segments are not homogeneous, which means there is likely to be a change point between the segments. Recursive merging follows a segmentation-based approach to change detection.
- **Predictive:** Predictive approaches to change detection are based on the assumption that one can learn a model for a portion of the input time series, and detect change based on deviation from the model. The underlying model can range from relatively simple smoothing models to more sophisticated filtering and state-space models. The change detection algorithm used to generate the Burned Area Product (a well-known MODIS data set) follows a predictive approach. This algorithm performs very poorly in parts of North America such as California [19] as illustrated in Figure 2. In addition, such products are geared towards specific kinds of changes (such as fires), and are not capable of detecting the broad set of changes can potentially be addressed (such as those due to deforestation, floods, droughts and insect infestations).

For a more comprehensive discussion of related work in land cover change, and the broader problem of time series change detection, we refer the reader to [4].

3. Algorithms for Land Cover Change

This section provides a brief description of the three time series change detection algorithms that are being evaluated in this study.

3.1. **Recursive Merging Algorithm.** Segmentation based algorithms operate under the assumption that given time series can be partitioned into homogeneous segments and boundaries between

Figure 2. This figure illustrates the poor coverage of the Burned Area product in California. The figure is a screen shot from Google Earth that shows the boundary of a fire near San Diego in 2003 (red line), and the pixels detected by the Burned Area product (circular markers).

the segments represent change points. There are two commonly used strategies to segment the time series [14]. A top-down strategy recursively partitions the time series till a stopping criteria is met. A bottom-up strategy on the other hand recursively merges smaller units . Existing techniques for segmentation ignore many key characteristics of the underlying ecosystem data such as seasonality and variability. Here we discuss the recursive merging algorithm [5] that follows a segmentation approach to the time series change detection problem and takes the characteristics of the ecosystem data into account.

The main idea behind the recursive merging algorithm is to exploit seasonality in order to distinguish between points that have had a land cover change and those that have not. In particular, if a given location has not had a land cover change, then we expect the seasonal cycles to look very similar going from one year to the next; if this is not the case, then based on the extent to which the seasons are different one can assign a change score to a land location. Recursive Merging follows a bottom-up strategy of merging annual segments that are consecutive in time and similar in value. A cost corresponding to each merge is defined as a notion of the distance between the segments. We use Manhattan distance in our implementation of the algorithm, although other distance measures can be used. One of the strengths of the Manhattan distance is that it takes the seasonality of the time series into account because it takes difference between the corresponding months. The key idea is that the algorithm will merge similar annual cycles and most likely the final merge would correspond to the change (if a change happened) and would have the highest cost of merging. In case the maximum cost of merging is low, it is likely that no change occurred in the time series.

The algorithm described above takes into account the seasonality of the data but not the variability. A high cost of merge in a highly variable time series is perhaps not as reliable indicator of change as a moderate score in a highly stable time series. In recursive merging algorithms the cost for the initial merges can be used as an indicator of the variability within each model. To account for this variability, the change score is defined as the ratio of the maximum merge cost (corresponding to difference in models) to the minimum merge cost (corresponding to the intra-model variability). Time series with a high natural variability, or time series with noise data due to inaccurate measurement have a high minimum cost of merging also, thus a smaller change score. As we show in Section 5.5 this method incorporates handling of noise and reduces false alarms in change detection. We will refer to this scheme as RM0.

3.2. **Lunetta et al. Scheme.** This anomaly based method for identifying changes relies on the fact that in a spatial neighborhood most of the locations remain unchanged and only a few locations get changed at any particular time interval.

For every location the algorithm computes the sequence of the annual sum of vegetation index for each year. The difference between the annual sum of consecutive years is then computed. We will refer to this as diff-sum. This is equivalent to applying first-order differencing [7] to the time series of annual sums. High values for the difference in the annual sum for consecutive years indicate a possible change. To determine the "strength" or "significance" of this change, Lunetta et al. compute a z-score for this diff-sum value for the combination of each year boundary and spatial location. When computing the z-score, Lunetta et al. define the standard deviation across *all* the spatial neighbors of the pixel for that time window in the data set and they further assume that diff-sum is normally distributed with a mean of 0 in the spatial neighborhood. An implicit assumption made by the scheme (due to this method of z-score computation) is that at each yearly boundary, same fraction of locations undergo land cover change. Note that high values of z-score indicate a decrease in vegetation and vice-versa. In subsequent discussions, we will refer to the scheme described above as the `LUNETTA0` scheme.

3.3. **CUSUM.** Statistical parameter change techniques assume that the data is produced by some generative mechanism. If the generative mechanism changed then the change will cause one of the parameters of the data distribution to change. Thus changes can be detected by monitoring the change in this parameter. CUSUM technique is a parameter change technique that uses the mean of the observations as a parameter to describe the distribution of the data values. The basic CUSUM scheme has an expected value μ for the process. It then compares the *deviation* of every observation to the expected value, and maintains a running *statistic* (the cumulative sum) CS of deviations from the expected value. If there is no change in the process, CS is expected to be approximately 0. Unusually high or low values of CS indicate a change. A large positive value if CS indicates an increase in the mean value of the vegetation (and vice-versa). We will refer to this scheme as CUSUM_MEAN.

4. Experimental Evaluation

4.1. **Earth Science Data.** The Earth Science data for our analysis consists of snapshots of measurement values for a vegetation-related variable collected for all land surfaces. The data observations come from NASA's Earth Observation System (EOS) [1] satellites and the data sets are distributed through the Land Processes Distributed Active Archive Center (LP DAAC) [2].

The specific vegetation-related variable used in this analysis was the enhanced vegetation index (EVI) product measured by the moderate resolution imaging spectroradiometer (MODIS) instrument (although any other vegetation index such as FPAR or NDVI could have been used). EVI is a vegetation index which essentially serves as a measure of the amount and "greenness" of vegetation at a particular location; Figure 3 shows a snapshot of EVI for the globe. MODIS algorithms have been used to generate the EVI index at 250-meter spatial resolution from February 2000 to the present; in this paper, the temporal coverage of the data is from the time period February 2000—January 2009.

4.2. **Evaluation Data Set.** Since our ground truth is about forest fires in California we created two data sets `DS1` and `DS3` which consists of forest pixels in California as described below.[1]

[1] A land cover map obtained from the Ecosystem Modeling Group at NASA Ames Research Center was used to subset forest pixels. The following land cover classifications were considered forest: Evergreen Needleleaf, Evergreen Broadleaf, Deciduous Needleleaf, Deciduous Broadleaf Forest, Mixed Forests.

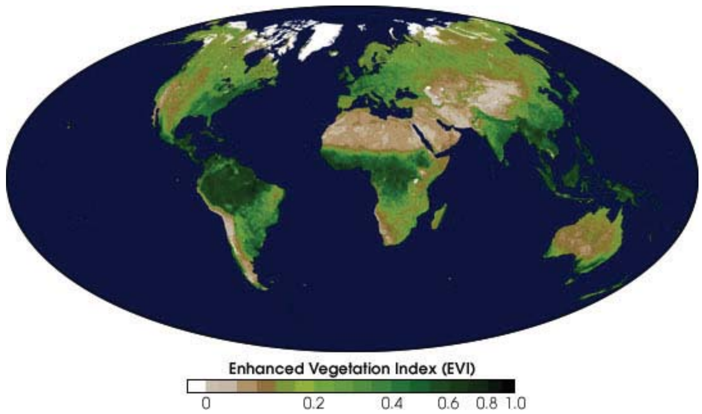

FIGURE 3. The above MODIS Enhanced Vegetation Index (EVI) map shows the density of plant growth over the entire globe for October 2000. Very low values of EVI (white and brown areas) correspond to barren areas of rock, sand, or snow. Moderate values (light greens) represent shrub and grassland, while high values indicate temperate and tropical rainforests (dark greens).

Source: MODIS Land Group, Alfredo Huete and Kamel Didan, University of Arizona.

`DS1`: **(Highest quality data).**
To create `DS1`, we preprocessed the data to eliminate poor-quality measurements by performing the following steps:

(1) The MODIS quality assurance (QA) flag (which describes atmospheric and sensor conditions under which the spectral measurements were taken) was used to retain only those observations of good quality, removing all observations that were tagged as being of *marginal* or of *low quality*. Another filtering step performed (recommended by earth science domain experts) was the removal of EVI measurements less than or equal to 0 and above 0.9.
(2) To reduce the impact of quality filtering, we converted the biweekly data to monthly data by averaging (using a simple mean) the available data for every month.
(3) We then discarded all locations that contained any missing data. In other words, the data for a location is retained only if the entire time series is available with no missing values and no low quality data.

`DS3`: **(Unfiltered data).**
`DS3` consists of the raw data without any processing for quality, i.e., the quality flag is not examined and we do not filter observations outside the recommended valid range.

The key characteristics and properties of the two data sets are summarized in Table 1. Note that by permitting noisy values, there is an over *five-fold* increase in the spatial coverage.

Data Set	# of pixels (N)	Frequency	Length of Time Series (T)	Noise	Missing Data
`DS1`	148,770	Monthly	108	Low level	No
`DS3`	787,777	Biweekly	207	High level	No

TABLE 1. Summary of evaluation data sets.

4.3. **Ground Truth Data.** Change detection studies are frequently plagued by the lack of good ground truth data [18] which forces the evaluation process to be more qualitative in nature. This

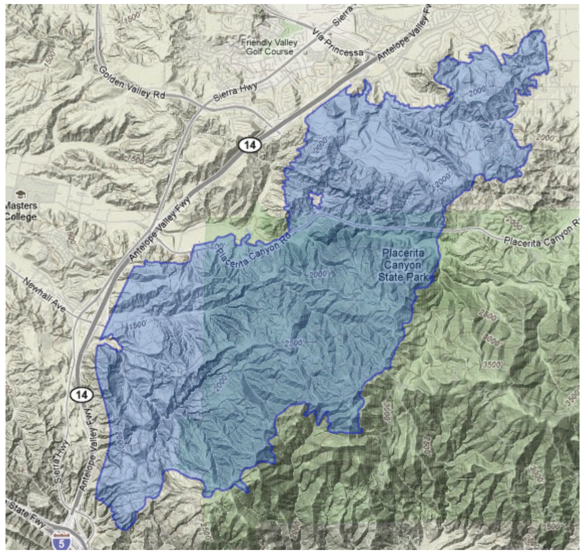

FIGURE 4. Example of a polygon representing the boundary of a fire.

frequently makes it difficult to objectively answer the question: *what is a change?*. In this study, we have utilized high quality ground truth data for fires generated by an independent source, and are thus able to perform an objective quantitative evaluation. We obtained fire boundaries generated by the state of California for the fire seasons for the years 2000 through 2008.

The ground truth data is in the form of *polygons* which represent the boundaries of forest fires. Each polygon $\mathcal{P}$ is a closed shape that consists of N sides (N is usually in the hundreds), with each vertex represented as a latitude/longitude coordinate pair, and may contain one or more holes $\mathcal{H}_i, i = 1 \ldots n$. The boundary of an individual fire is then $\mathcal{P} \setminus \{\mathcal{H}_1 \cup \mathcal{H}_2 \cup \ldots \cup \mathcal{H}_n\}$. A example of a fire boundary is shown in Figure 4; the fire occurred in 2004 near Santa Clarita, CA. The blue filled region represents the polygon; the dark blue line is the outside boundary of the polygon, while a hole can be seen in the middle of the map.

There are two issues with the ground truth that we are using for evaluation. First, there are changes in California forests due to other reasons that fire(e.g due to logging). Since they are not part of the ground truth, they will be considered false positives if they are discovered by the change detection algorithm. Second issue arises due to the inaccuracy of the forest filter due to which many non-forest locations such as farms also become part of our data set. These locations may have actual change that is detected by the algorithm but again it will appear to us as false positive. However, we expect these issues will impact all the algorithms similarly and thus we will still be able to make judgement about there relative performance.

4.4. **Evaluation Methodology.** Given a time series data set D with N pixels, let us assume that any 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 ground truth data set which consists of the true labels of each of the pixels; let M be the *total* number of actual disturbances as determined by ground truth. To evaluate the performance of a given change detection algorithm at rank n, we count the number of true disturbances in the top n portion of the sorted change scores of all the pixels, where n is the number of actual disturbances ($1 \leq n \leq M$). 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. Specifically, performance is measured in terms of the number of instances correctly identified and the number of instances

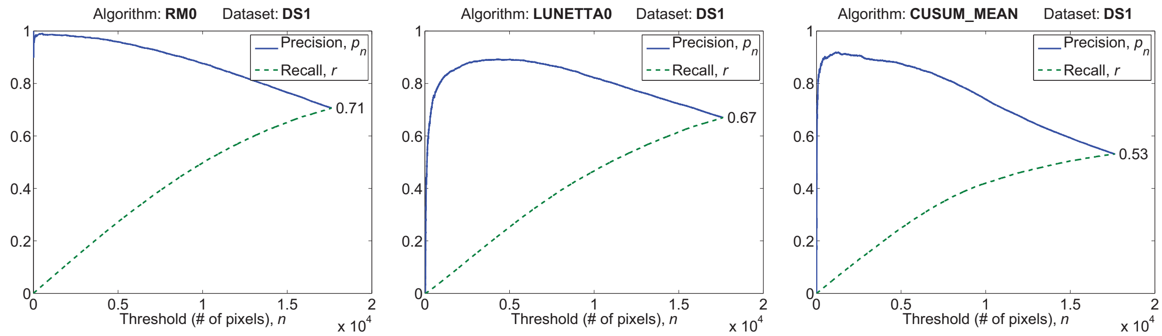

FIGURE 5. Comparison of algorithms on DS1.

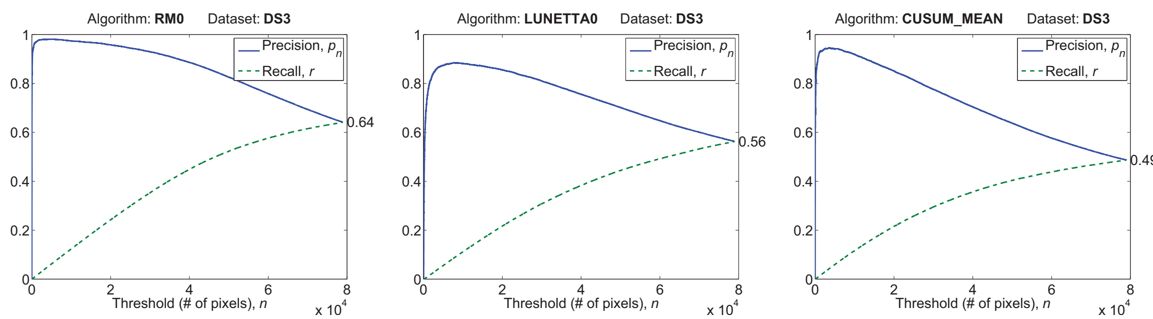

FIGURE 6. Comparison of algorithms with noisy data (DS3).

missed in the top-n ranked instances. We use a precision metric (called p_n) employed in context of information retrieval [3] and anomaly/outlier detection [6], which is appropriate for the top-n ranked setting. The performance metrics are defined as follows:

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

Note that as n increase, p_n will tend to decrease (a greater fraction of lower scoring points are likely to be false positives) and recall will increase (since, eventually for large enough n, all true positives will be included). One specific value of interest is the one when n is equal to the number of fire pixels (ground truth). At this value of n, $p_n = r_n$ since $TP_n + FP_n =$ M. Also, if the change detection algorithm does the perfect job of identifying fires, then upto this value of n, p_n will remain 1 (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).

4.5. **Experimental Results.** The three algorithm were run on datasets DS1 and DS3. Figure 5 and 6 shows precision and recall curve for each algorithm as n changes from 1 to the number of fire pixels in the ground truth for each dataset (18450 in DS1 and 82311 in DS3). Tables 2 and 3 show overall results (aggregate count) broken down by each year. It is to be noted that the false positives labelled by the ground truth can either be time series incorrectly classified as change by the algorithms or can be changes other than fires like logging, conversion to golf course etc.

4.6. **Observations.**

(1) **Performance is better on DS1 than DS3**
Figure 5 and 6 show that all the three algorithms perform better on DS1 than DS3. The primary reason is that data set DS3 has no quality filtering and thus contains time series

Year	# of pixels in fire polygons Polygon Size	RM0	LUNETTA0	CUSUM_MEAN
2000	111	54	39	12
2001	1142	814	850	1009
2002	2407	1383	2119	2164
2003	4946	3609	3670	4338
2004	661	423	463	521
2005	192	96	128	134
2006	278	146	197	152
2007	1935	1413	1353	1360
2008	6778	5312	3811	490
SUM	18450	13250	12630	10180
$p_n(= r_n)$	1.00	0.72	0.68	0.55

[1] The second column shows the # of pixels in the data set that fall in the fire polygons
[2] The next three columns show the number of pixels detected by respective change detection algorithms that fall in the fire polygons.

Table 2. Results of algorithms on DS1.

Year	# of pixels in fire polygons Polygon Size	RM0	LUNETTA0	CUSUM_MEAN
2000	1379	458	58	443
2001	6827	3661	105	5520
2002	12114	7061	1238	9335
2003	12292	8514	937	9915
2004	4218	2786	857	3152
2005	744	293	115	336
2006	6165	3900	442	3948
2007	10671	9285	423	6736
2008	27901	17581	2423	1742
SUM	82311	53539	6598	41127
$p_n(= r_n)$	1.00	0.65	0.08	0.50

[1] The second column shows the # of pixels in the data set that fall in the fire polygons
[2] The next three columns show the number of pixels detected by respective change detection algorithms that fall in the fire polygons.

Table 3. Results of algorithms on DS3.

which are highly noisy. These time series can receive artificially high change score due to noisy values.

(2) **RM0 outperforms LUNETTA0 and CUSUM_MEAN**
Figure 5 and 6 show that RM0 consistently performs better than LUNETTA0 and CUSUM_MEAN on both the datasets DS1 and DS3. The difference in performance is especially significant on the dataset DS3; The reason is that DS3 has more time series that are highly variable because of no quality filtering and RMO is able to perform better since it has a built-in notion of variability modeling (we illustrate this is greater detail in the next paragraph). The following illustrative examples highlight the difference between the three algorithms in

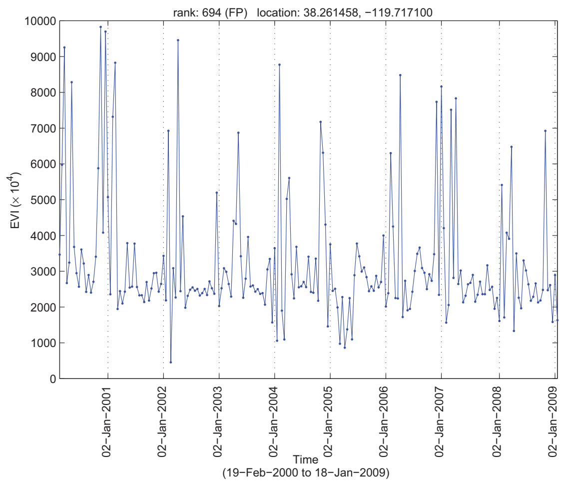

FIGURE 7. Sample of a false positive detected by CUSUM_MEAN_MISSING on DS3.

their ability to handle variability in time series. Figure 7 shows a false positive that was detected by CUSUM_MEAN but *not* by RM0 and Figure 8 shows a false positive that was detected by LUNETTA0 but *not* by RM0. RM0 due to its ability to account for variability gives these time series a low change score and does not detect them as change points.

(3) **RM0 does well because it takes into account the variability in the time series**
To assess the ability of RM0 to model variability, we evaluate a variation of RM0 that does not perform the normalization step of RM0 (i.e we do not divide the score by the minimum of the scores of merging). We refer to this scheme as RM_NO_NORM. Figure 9 shows the precision and recall curve for the RM_NO_NORM. It can be observed that the performance of RM_NO_NORM degrades severely compared to RM0 especially on the dataset DS3. As an illustration, Figure 10 shows a time series that is given a high change score by RM_NO_NORM but not by RM0.

(4) **LUNETTA0 can be improved by eliminating normalization**
Table 4 and Figure 11 show the number of pixels burned in each year from 2001 to 2008 on the DS1 data set. Also shown is the standard deviation of the annual differences corresponding to each year. The data indicates that the standard deviation of annual differences is higher for time periods when a greater number of pixels are burned (Similar conclusions were drawn for DS3). For these years (especially 2008), the change scores will be diminished compared to a year such as 2006. This means that if pixel n_i has a fire in 2006 and pixel n_j has a fire in 2008 and they have *exactly* the same time series, pixel n_i will receive a higher change score than pixel n_j. Thus, we observe that the normalization step performed in Lunetta can lead to a suboptimal change score when there is a difference in the variability of delta over different years (which is what happens in the case of forest fires). To test this observation, we implemented a variation of LUNETTA algorithm that skips the normalization step. We refer to this scheme as LUNETTA_NO_NORM. From Figure 12 it is clear that LUNETTA_NO_NORM performs better than the original Lunetta scheme. However, it is to be noted that LUNETTA_NO_NORM still performs worse than RM0.

5. Conclusion

A number of insights can be derived from the quantitative evaluation of the algorithms and their variations presented in this paper. On relatively high quality datasets, all three schemes perform reasonably well, but their ability to handle noise and natural variability in the vegetation data

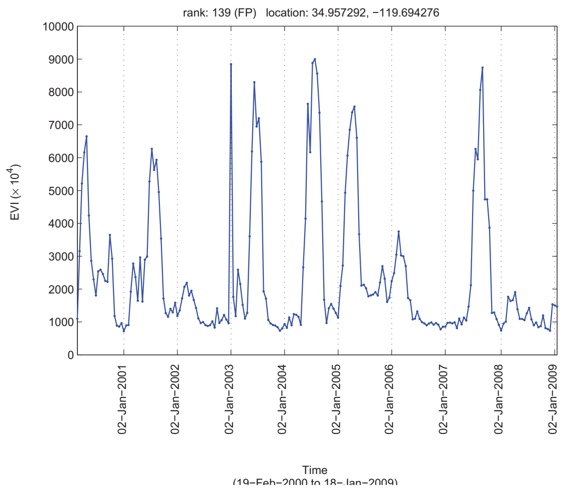

FIGURE 8. Sample of a false positive detected by LUNETTA0 on DS3.

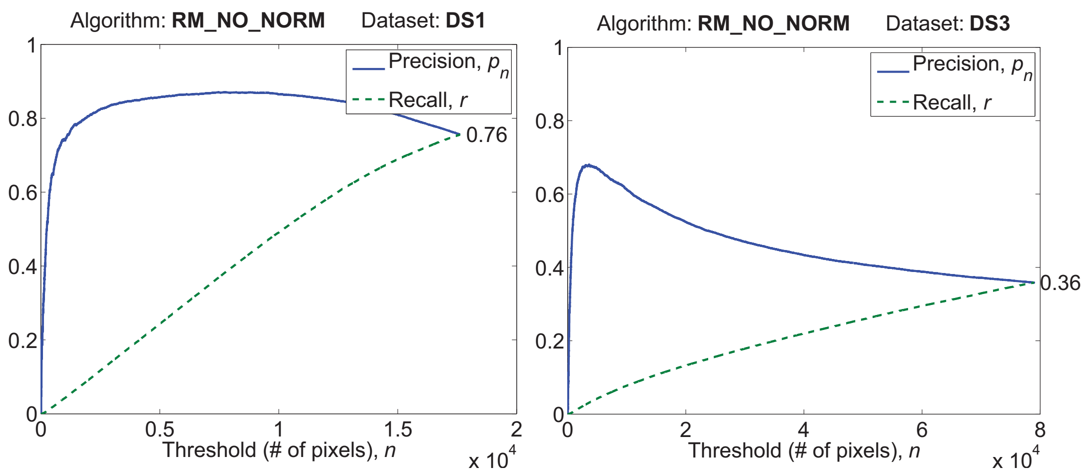

FIGURE 9. Performance of RM_NO_NORM on DS1 and DS3.

differs dramatically. In particular, Recursive Merging algorithm significantly outperforms the other two algorithms since it accounts for variability in the time series.

However, the algorithm has several limitations that need to be addressed in future work. For example, due to manner in which the segments are constructed from annual cycles, changes occurring in the middle of segment boundaries are given lower scores than changes occurring at the segment boundaries. The algorithm normalizes the change score for a given time series by the estimated variability. The normalization is currently performed using the minimum distance between a pair of segments, which is not optimal: Figure 13 illustrates how this normalization leads to false positives when a time series with relatively low mean undergoes a small shift.

Additionally, there are several limitations of the experimental evaluation in this study. For example, the ground truth data set consists of only one type of land cover change (forest fires), thus excluding many other changes of interest. Furthermore, the nature of vegetation data in California can be quite different from other parts of the world such as the tropics, where the issues of noise are acute because of persistent cloud cover.

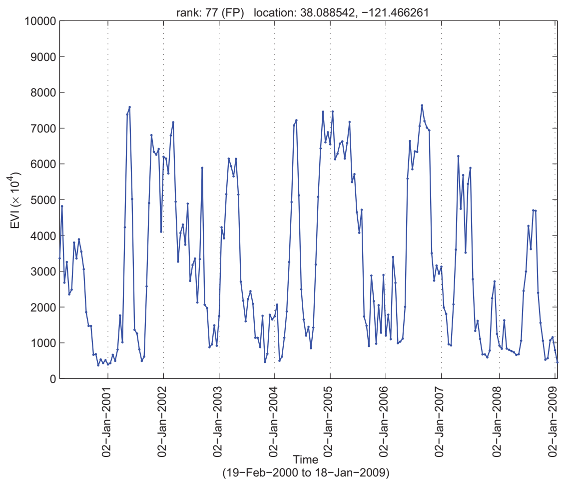

FIGURE 10. Sample of a false positive detected by `RM0_NO_NORM` on `DS3`.

Year	# of fire pixels	Standard Deviation
2001	1142	0.20
2002	2407	0.32
2003	4946	0.27
2004	661	0.27
2005	192	0.28
2006	278	0.21
2007	1935	0.29
2008	6778	0.37

TABLE 4. Standard deviation of integrated annual differences on `DS1`.

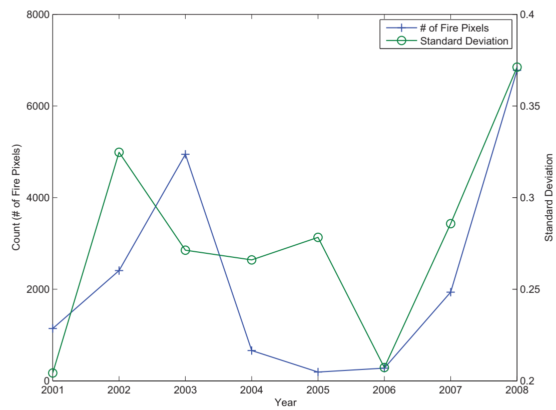

FIGURE 11. A visual representation of the data in Table 4.

ACKNOWLEDGEMENTS

We are grateful to the anonymous reviewers for their comments and suggestions, which improved this paper. This work was 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 would also like to thank Ivan Brugere, Divya Alla, Vikrant Krishna and Matt Kappel for their helpful comments.

REFERENCES

[1] NASA Earth Observing System.
`http://eospso.gsfc.nasa.gov`.

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

[3] R. Baeza-Yates and B. Ribeiro-Neto. *Modern Information Retrieval.* Addison Wesley, Reading, MA, 1999.

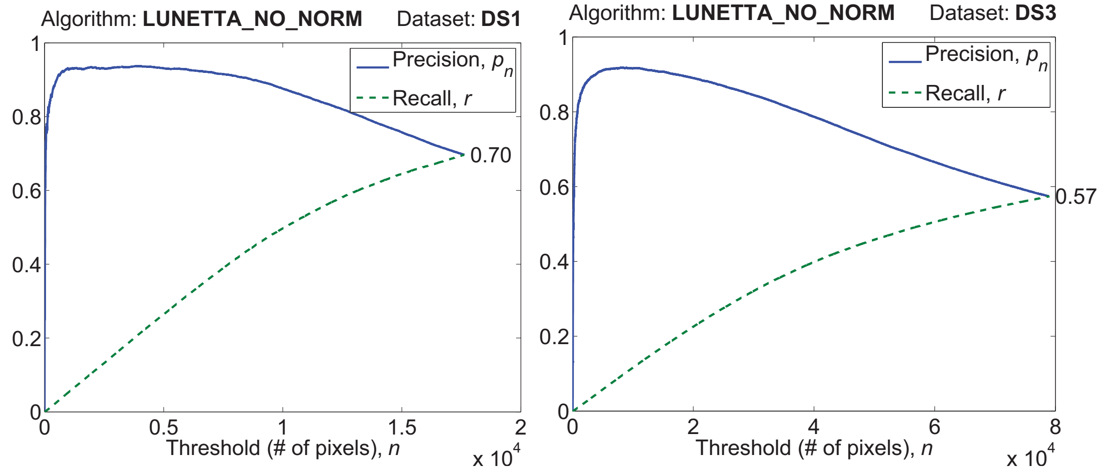

FIGURE 12. Performance of LUNETTA_NO_NORM on DS1 and DS3.

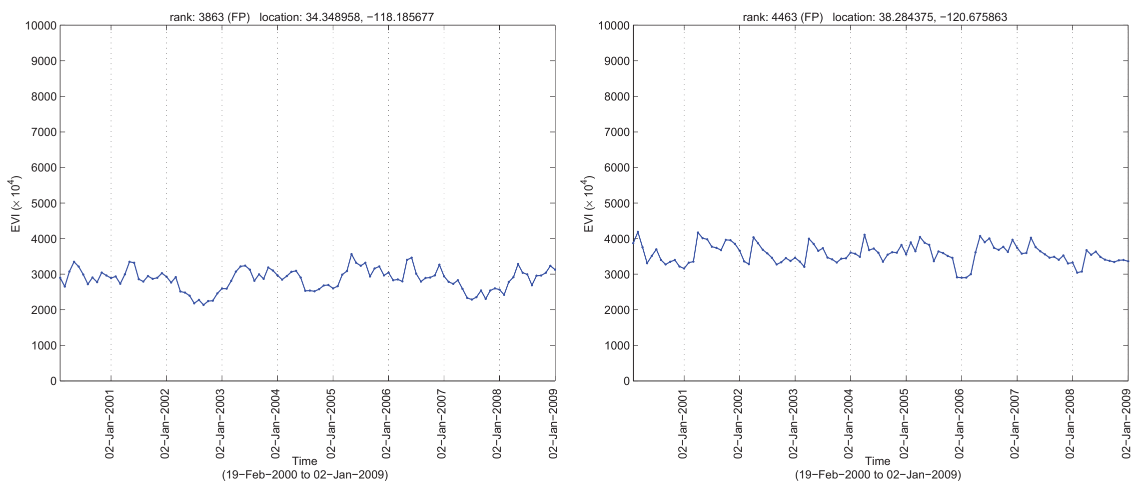

FIGURE 13. Sample of false positives detected by RM0 in DS1.

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

[5] S. Boriah, V. Kumar, M. Steinbach, C. Potter, and S. Klooster. Land cover change detection: A case study. In *KDD '08: Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 857–865, 2008.

[6] V. Chandola. *Anomaly Detection for Symbolic Sequences and Time Series Data*. PhD thesis, University of Minnesota, 2009.

[7] C. Chatfield. *The Analysis of Time Series: An Introduction*. Chapman & Hall/CRC, 2004.

[8] W. Cohen and M. Fiorella. Comparison of methods for detecting conifer forest change with thematic mapper imagery. In *Remote sensing change detection: environmental monitoring methods and applications. Chelsea,(MI): Sleeping Bear Press. p*, pages 89–102. Ann Arbor Press, Chelsea, MI, 1998.