INCORPORATING NATURAL VARIATION OF TIME SERIES IN THE CHANGE DETECTION FRAMEWORK TO IDENTIFY ABRUPT FOREST DISTURBANCES

VARUN MITHAL*, ASHISH GARG*, IVAN BRUGERE*, SHYAM BORIAH*, VIPIN KUMAR*, MICHAEL STEINBACH*, CHRISTOPHER POTTER**, AND STEVEN KLOOSTER**

ABSTRACT. Ability to monitor forest related change events like forest fires, deforestation for agriculture intensification, and logging is necessary for effective forest management. Time series remote sensing data sets such as Enhanced Vegetation Index (EVI) can be used to identify these changes. Existing approaches work on small data sets spanning over a specific geographic region of a homogeneous vegetation type. Also, most of these need training samples or require setting of parameters for each geographic region individually. These limitations make the algorithms unscalable and restrict their global applicability. In this paper, we present a scalable time series based change detection framework that overcomes these limitations of the existing methods. We introduce the concept of natural variation in EVI response of location and incorporate it in change detection paradigm. We evaluate the change events identified by our approach using forest fire validation data in California and Canada. The results of this study demonstrate that the inclusion of a measure of natural variability of a location in determining its change score improves detection accuracy, and makes the paradigm more robust across vegetation types and regions.

1. INTRODUCTION

Forests act as sink of atmospheric carbon and forest disturbances such as fires and deforestation cause the stored carbon to be released into the atmosphere. In addition, forests are the home to many ecosystems and these disturbances cause them severe damage. For efficient and effective management of forest resources, reliable and quantifiable observation of forest cover changes at a global scale is critical [15]. Some nations have allocated resources to monitor disturbances in their forests. As an example, Brazil has developed a system for deforestation monitoring called PRODES. Regional products such as these are infrequent because they require considerable monetary resources. Therefore, there is a need to develop a forest monitoring system to identify global forest disturbances.

Data collected from remote sensing instruments can be used to identify changes in forest. The bulk of work in identifying land cover changes using remote sensing data involves image comparison methods [6, 13]. These methods include classifying locations using the reflectance data and using post-classification comparison to identify changes. However, they have limited applicability because they require training data sets for supervised classification process. Due to the supervised learning step for classification involved, these techniques are region-specific and not globally scalable. Furthermore, the error in classification gets compounded during change detection. In addition, several characteristics like rate of change and exact change date cannot be found using these image comparison based methods because these approaches compare snapshots between two dates and information between those dates is not considered. Moreover, only time series approaches provide fine grained information about land cover dynamics that is necessary to quantitatively assess the carbon impact of land cover changes [17]. Publicly available global time series data sets such as Enhanced Vegetation Index (EVI) can be used to identify changes in forest cover. Hence there is an increasing interest in time series-based approaches to change detection in vegetation data [2, 4, 8, 10, 11, 12, 14, 18]. However, these methods have been used on a small data sets spanning over a local geographical area

*University of Minnesota, <mithal,ashish,ivan,sboriah,kumar,steinbac>@cs.umn.edu
**NASA Ames Research Center, chris.potter@nasa.gov, sklooster@gaia.arc.nasa.gov.

and comprising of a homogeneuous vegetation. These existing algorithms have a parameter setting step that is fine-tuned for performance in that specific geographical region and vegetation. This is a serious limitation that makes it difficult to apply the algorithms on a global scale.

There are primarily three types of time series-based land cover change detection approaches. Temporal segmentation methods divide the time series into homogeneuos regions and the boundary of the segments indicate the change in vegetation phenology [4, 10]. These approaches are aimed to identify any change in the vegetation phenology such as change from one land cover to other or changes in crops. The other approach is to look for trends in the time series spanning over multiple years to identify gradual decrease in vegetation response [8, 20]. Such gradual changes represent forest degradation such as due to beetle infestation, long-term droughts, etc. The third approach, which is the focus of the paper, is to identify an abrupt decrease in vegetation by predicting EVI values from a learnt model for the vegetation time series and using prediction error to identify a change [3, 8, 12, 14, 18]. Roy et al. [18] use a model-based prediction scheme for identifying fires from time series data and use this for generating the global Burned Area Product. Kucera et al. [12] use a CUSUM based technique and model fitting to identify forest fires in Portugal. Hammer et al. [8] use regression based technique to model the short and long term trends in NDVI for detecting deforestation in pan-tropical rain forests. Lunetta et al. [14] use a spatial anomaly detection method for identifying deforestation. Boriah [3] describe Yearly Delta, a model based approach that uses mean annual EVI difference between successive years to identify changes such as forest fires and show that is is comparable in performance to Recursive Merging proposed in [4] which was shown in [5] to significantly outperform algorithm based on CUSUM and the scheme proposed by Lunetta et al. [14].

A global scale analysis reveals that some vegetation types are highly stable and show a small decrease in EVI when a disturbance occurs, but this decrease can be significant compared to the otherwise stable nature of past EVI values. On the other hand, some vegetation types show random fluctuations in the EVI signal which occur due to atmospheric interference and various natural sources ranging from soil conditions to inter-annual temperature and precipitation. Thus, significance of Yearly Delta score differs based on the region and vegetation type, and the score thresholds need to be adjusted separately for different geographic regions and land cover types to avoid too many false alarms. In this paper, we propose two *time series change detection algorithms* that utilize the temporal structure present in the remote sensing data to incorporate natural variability in the vegetation response of a location in the Yearly Delta algorithm. The incorporation of the concept of natural variation in the Yearly Delta algorithm improves the change detection accuracy, and makes the paradigm more robust across vegetation types and regions. The evaluation of the proposed method is done quantitatively using validation data on forest fires in California state (USA) and Yukon state (Canada). We also compare performance of our algorithms against the output from Burned Area Product, a well-known global NASA product for fire monitoring by Roy et al. [18]. The evaluation results assert the importance of incorporating natural variability of vegetation response in change detection. In particular, the experiments illustrate the need for variability modeling if change detection is performed on larger regions with multiple vegetation types.

1.1. **Key contributions.** The key contributions of this paper are: (1) a scalable framework to identify significant abrupt changes in spatial-temporal remote sensing vegetation data sets to address the problem of land cover change detection, (2) introducing a concept of natural variation of EVI response in the identification of changes in EVI signal, (3) a method to associate significance to observed annual changes in EVI with respect to the natural variability of the location, and (4) quantitative evaluation of the performance of the proposed approach using validation data sets available for forest fires in California (USA) and Yukon (Canada) and also comparison with an existing well-known global fire monitoring product.

1.2. **Organization of the paper.** In the paper, we describe the data used in Section 2. Section 3 presents the proposed change detection framework and the details of the proposed algorithms.

Section 4 describes the validation data and evaluation methodology for this paper. Section 5 provides analysis of the results and Section 6 discusses the key challenges in the task of land cover monitoring, limitations of the proposed algorithms and the future research directions.

2. Data and Preprocessing

Global remote sensing data sets are available from a variety of instruments at different spatial resolutions as a sequence of global snapshot of measurement values. In principle, the proposed algorithms can be applied to any geospatial dataset that features regular, repeated observations, consistent image registration and well-defined composite indicators of vegetation. In this study, we use the Enhanced Vegetation Index (EVI), a data product (MOD13Q1) derived from measurements taken by the Moderate Resolution Imaging Spectroradiometer (MODIS) sensor on NASA's Terra satellite and distributed through the Land Processes Distributed Active Archive Center [1]. EVI essentially measures the "greenness" signal (area-averaged canopy photosynthetic capacity) as a proxy for the amount of vegetation at a particular location. MODIS data has been used to generate a continuous record of the EVI index at spatial resolution of 250 meters from February 2000 to the present. This index is generated at a temporal frequency of 16 days: each instance in the product is composited using the highest quality data from 16 daily raw observations.

In this study we use MODIS EVI data for California (USA) and Yukon (Canada). The data of California is at 250 m spatial resolution and 16 day temporal resolution. A spatial mask of MODIS landcover category [7] was used to separate land cover categories of interest (forests, savannas and shrubs) from other categories like agriculture and urban using the MODIS product for land cover maps (MCD12Q1). The data set *DSCalifornia* has 3,389,564 pixels and predominant changes include forest fires, deforestation and urban expansion. The other dataset *DSCanada* used in this study is MODIS EVI at 1km spatial resolution for Yukon Province in Canada. The data set has 551,275 pixels with the major change type as forest fires. This data set has significant homogeneity and no MODIS land cover mask was used. Winter months for this data set were pre-processed to an EVI value of 0 because winters are snow clad at this latitude and have no vegetation response. The EVI values are in the range 0 to 1, but in this study we scale the values by a factor of 10,000 and so the data values range from 0 to 10,000.

3. Algorithm Description

In this section, we describe the two proposed change detection algorithms to identify abrupt decrease in vegetation such as due to forest fires and deforestation. The main intuition behind these algorithms is that in stable forests EVI values for future time steps tend to be similar to previous years while accounting for seasonal variation. On the other hand, changes like fires and deforestation are characterized by an abrupt decrease in EVI after the change. The algorithms build a model used for predicting the expected EVI values for the future years that is deviation of the future observations from this model indicates a change. A measure that quantifies the deviation of future observations from the model is used to assign the change score. One possibility is to predict the EVI for each time step, and use the prediction error as change score. However, this scheme is susceptible to noise in the data and causes too many false alarms. One approach to address this issue of noise susceptibility is to use the notion of persistent decrease, and flag a time series for change only if a significant fraction of time steps from a time window of size m exceed the threshold. This idea of persistence is used by Roy et al. [18] in their algorithm to identify fires from daily remote sensing time series. Another approach is to predict a more stable statistic, for example, the mean of the successive time steps over an year that is more robust to noise than deviation from prediction at a single time step. This approach is used in the Yearly Delta algorithm discussed in [3] and [15].

In the following, we describe the Yearly Delta algorithm and its two new variations: Variability-Aware Yearly Delta and Vegetation-Independent Yearly Delta. Our focus is to incorporate the natural variation of EVI time series for a location in assigning the change score. This concept of variability proposed in the paper is applicable to both approaches, Burned Area and Yearly Delta.

3.1. **Yearly Delta algorithm (YD).** The main intuition behind this algorithm originally presented in Boriah [3] and also used in Mithal et al. [15] is that for the time step corresponding to the date of abrupt change event, the difference between the annual mean EVI of the previous and following year will be high. YD algorithm considers the previous year as the model and assigns a change score to each time step as the difference between the mean annual EVI of the previous year and the following year. The YD score for a location is the maximum change score across all time steps and the time step with the maximum score is considered the time point for the change. The pixels are ranked based on their YD score and a certain number of the top ranked pixels are considered as changes. This algorithm works under the assumption that the pixels which have an abrupt change event will get a higher score than those which do not have an abrupt change event because the undisturbed pixel typically do not have an unusually large EVI decrease from one year to next. Figure 1(a) shows a location in California with a fire occurrence. We can see an abrupt decrease in EVI value after the fire in year 2007 that will lead to a high mean annual EVI difference and therefore a high YD score.

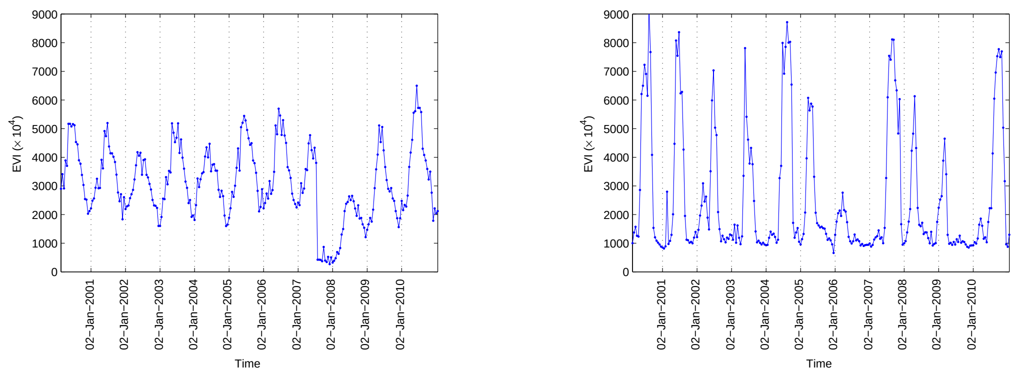

(a) EVI time series for a pixel in forest location with YD score of 2500 and VD score of 2100 for year 2007.

(b) EVI time series for a pixel in shrub location with YD score of 2500 and VD score of 700 for year 2006.

FIGURE 1. Illustrative examples to understand the power and limitations of the YD algorithm

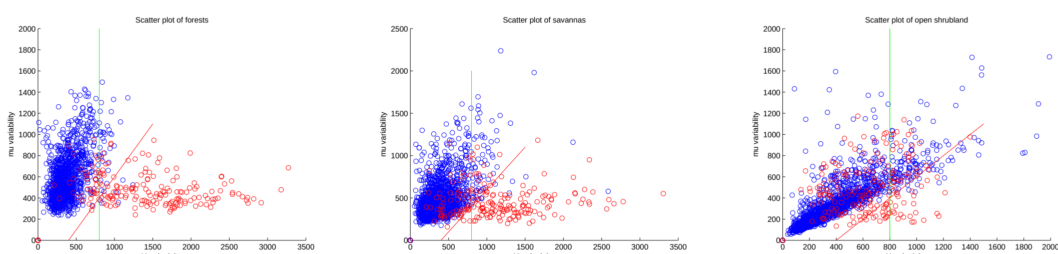

(a) MODIS Forests

(b) MODIS Savannas

(c) MODIS Shrubs

FIGURE 2. Scatter plot of mean variability (μ_{var}) against the YD score for different land cover categories in California.

3.2. **Variability-Aware Yearly Delta (VD).** The observations in the future years vary from the model built based on previous years due to natural variability arising from changes in weather, soil conditions, etc. Shrub land cover shows a variety in the natural variation in their EVI signal and are severely affected by changes in local climate conditions of precipitation and temperature as opposed to forests which are more resilient to such climatic changes. Thus, the mean annual EVI differences are not expected to be always 0 for the unchanged locations. Also, this variability depends on the location's local geography and land cover type. The change score should reflect the significance of the deviation with respect to the natural variation of vegetation response for that location. The YD algorithm does not make use of information about the natural variation in EVI response. The same YD score can correspond to a significant loss of vegetation for some vegetation type or can occur due to natural variation in others. To understand this consider the two time series in Figure 1(a) and 1(b) that have the same YD score. Figure 1(a) corresponds to an actual change in year 2007, while Figure 1(b) has the same YD score for year 2006 due to inter-annual natural variability that exist in shrubs due to high sensitivity to changes in climate variables such as precipitation. The limitation of YD to distinguish between the two types of changes illustrated by Figure 1(a) and 1(b) motivates the need to incorporate the notion of natural variation in the change detection paradigm. One possibility is to use the differences in EVI values in the past to model the natural variation of a location. In this scheme, each annual segment in first k years is considered a model and the remaining $k-1$ segments are considered observed values and the mean Manhattan distance for each of the pairs is computed to give a distribution of variability scores for that location. A YD score that lies in this distribution is likely to occur even by random fluctuation. So we modify the score as YD score relative to the mean of this distribution (μ_{var}) for each location. The new score is called the VD score and is computed as VD score = YD score - μ_{var}

To illustrate the advantage of subtracting the μ_{var} from YD score, we show the scatter plots of YD score against μ_{var} for a random sample derived from three different land cover categories in California. Figure 2 shows the unchanged locations as blue circles and changed locations as red circles. The vertical line in green shows the constant YD score of 800 and the oblique line in red shows the constant VD score of 400. These scores were chosen for the two algorithms because they gave similar number of changed events. Circles lying to the right half of these lines will have change scores higher than for that line and will be detected as changes by the algorithms respectively. So the blue circles in this right half are the errors of the algorithms. We see in Figure 2(a) that both algorithms will correctly identify the fires on the forest land cover though after subtracting variability we will reduce some of the errors. Figure 2(b) shows the same plot for Savannas. Here we notice that YD will make more errors as compared to VD and incorrectly label a few unchanged locations as changed. The scatter plot for Open Shrublands is shown in Figure 2(c) and we see that this is a difficult category and performance of both YD and VD is poor. However, the number of mistakes made by VD is significantly lower to those by YD. These scatter plots illustrate the utility of modeling variability in the change score especially in the highly variable land cover types. Similarly, in the Figures 1(a) and 1(b), we see that the two locations get the same YD score, but VD is able to incorporate the inherent variability of the locations and gives very different change scores.

Our experiments show that any value for k between 3 and 5 works well. Since the first k years are used for modeling natural variability, the change detection starts from $k+1$ year and changes in the first k years are not detected. Also note that this method assumes that the first few years that are used for variability modeling are undisturbed in the location. If a change event occurs during these initial years, it will cause the location to get a high variability score and a later change at that location will go undetected. This limitation can be addressed by using the previous k years instead of the first k years for computing variability under the assumption that abrupt changes such as fires and deforestation do not happen multiple times in k years.

The VD algorithm highlights the importance of incorporating the natural variation of the vegetation response at a location in computation of the change score. In the discussion above, we see that μ_{var} is a good indicator of expected natural variations in mean annual EVI differences and using the VD score can significantly improve the performance, especially in some land cover types. Boriah et al. [5] illustrated a similar advantage of modeling variability in the context of their segmentation algorithm Recursive Merging.

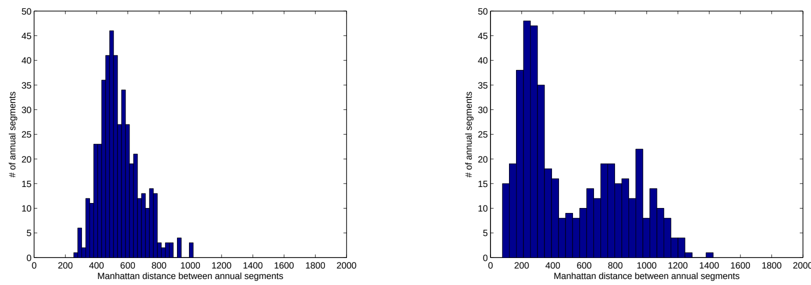

(a) The distribution of mean Manhattan distances between annual segments for a cluster of locations with smaller spread.

(b) The distribution of mean Manhattan distances between annual segments for a cluster of locations with a wider spread.

FIGURE 3. The distribution of variability scores (mean Manhattan distances between annual segments) of groups of pixels from two different locations in California with same μ_{var} (i.e. around 500).

3.3. **Vegetation-Independent Yearly Delta (VID).** Consider the scenario where the distribution of the mean Manhattan distances between annual segments for two locations in different types of vegetation have the same mean (μ_{var}) but different spread. In this scenario, if the same decrease in mean annual EVI was noticed in the two locations, the VD algorithm will give them the same score. However, the probability that the decrease in mean annual EVI was observed by a random chance is different for the two locations. For the location with smaller spread of distribution (i.e. smaller standard deviation) the probability that this mean annual difference is by random chance is lower, and for the location with a wider spread (i.e. higher standard deviation) this probability is higher. VD, which will assign same score, is unable to distinguish between the two cases. This is a serious limitation for VD if it is used for a composite data set with multiple types of vegetation. In such a scenario, locations will have different spread of their variability score distribution and a higher VD score threshold will miss the actual change that occurred in the more stable vegetation and therefore have a poor recall for those vegetation types. On the other hand, a lower VD score threshold will have many false positives from locations that have a wider spread of variability distribution because unchanged pixels will also have same VD score by random chance.

As an illustrative example, two locations in *DSCalifornia* were chosen and the 30 nearest neighbors for two locations were computed based on Manhattan distance measure in their EVI time series. The mean Manhattan distance between annual segments for each pixel in the two groups were computed and Figure 3 shows the distribution of the mean manhattan distances between annual segments of the two groups. These groups have similar mean variability (μ_{var}) but different spread in the distribution of the variability scores. The same mean annual EVI decrease observed has a different probability of occurring by natural variation in the two vegetation types. For example, if the YD score was 1000 then the probability that this would be observed by natural variation in 3(a) is considerably small, but 3(b) has a high probability of getting this score by natural variability. Thus, there is a need to further scale the VD score with the standard deviation of the variability score distribution to accurately estimate the significance of the change.

The VID algorithm tries to address this limitation by including the standard deviation of the variability in the change score. It assumes that the random fluctuations in mean annual EVI for a particular vegetation type are normally distributed for a location and estimates the mean μ_{var} and standard deviation σ_{var} of the variability score distribution as the maximum likelihood estimates

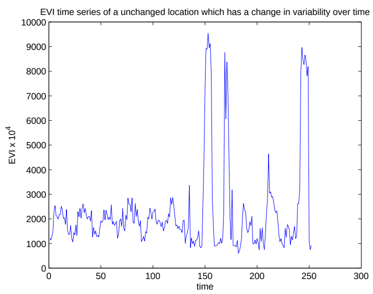

FIGURE 4. EVI time series for a location in California with highly stable initial years (before time step 100) for which the variability modeling was done and larger variations due to climatic variability in later years. Such locations are incorrectly identified as land cover change by the VID algorithm.

for the distribution. The new score is called the VID score and is computed as VID score = (YD score - μ_{var}) / σ_{var}.

This score can be viewed as the z-statistic from the standard normal distribution. A high VID score threshold implies a lower false positive rate and vice-versa. In addition, fixing the same VID score threshold for all locations will incur same false positive rate across vegetation type. This, however, depends on the assumption that for different vegetation types the variability scores have a near normal distribution and the futute EVI values also follow the same distribution if there is no change event. We observe that the assumption is true in most cases and the false positive rate for the algorithm is independent of the vegetation type. Due to climatic and other factors, for some vegetation types this assumption is not true and the false positive rates are higher for these vegetation types for the same score threshold. As an example, see Figure 4 which is the EVI time series of a location in California in which the variability changes with time perhaps due to changes in precipitation.

Note that we add a small number (1% of EVI scale) to the estimate value of σ_{var}. In case the σ_{var} is close to 0 for a very stable location, this avoids a small change from getting an extremely high VID score. This is especially important for locations with highly stable EVI such as in arid and semi-arid areas, where slightly high vegetation response for an year due to higher precepitation might lead to a false alarm due to a high VID score.

4. Evaluation

We use the same evaluation strategy as described in Boriah et al. [5] to understand the relative performance of different change detection techniques. The following describes the validation data used in this study and provides a brief overview of the evaluation methodology.

4.1. **Validation Data.**

Change detection studies are frequently plagued by the lack of good ground truth data [16] which forces the evaluation process to be more qualitative in nature. In this study, we have utilized high quality validation data for fires generated by an independent source, and are thus able to perform an objective quantitative evaluation. Specifically, we obtained fire boundaries generated by the state of California for the fire seasons for the years 2006 through 2008 and the state of Yukon in Canada for years 2004 to 2008. The validation data is in the form of *polygons* which represent the boundaries of forest fires. Our EVI data is georeferenced by the latitude and longitude value for the pixel center. Thus, a pixel is considered inside a polygon if the pixel center is inside it otherwise it is considered outside the polygon.

The histogram in Figure 5(a) shows the distribution of land cover type of the pixels that lie *inside* the validation polygons for California; i.e., these are the pixels which actually burned according to the validation data. The figure shows that shrubland and savannas account for a significant portion

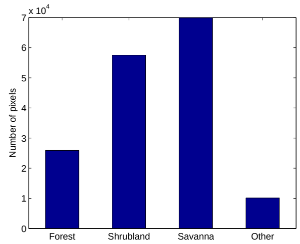

(a) Landcover distribution inside the fire polygons of California.

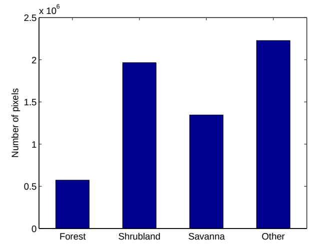

(b) Landcover distribution in California

FIGURE 5. The histograms show the number of pixels of each land cover type (using the MODIS land cover map) inside the fire polygons and in entire California.

		Predicted	
		Fire	No Fire
Validation Data	Fire	TP_n	FN_n
	No Fire	FP_n	TN_n

TABLE 1. Confusion matrix.

of the burned regions in California. The land cover types included in our study are MODIS forests, savannas and shrubs and we exclude the pixels belonging to the "other" MODIS landcover category in our California data set. This is because land cover categories such as agriculture that belong to the "other" category have a vast number of changes and majority of these changes are not related to fires and therefore will be considered as false positives by our validation data. Also, the fraction of pixels belonging to the "other" category is the ground truth is small compared to their fraction in the entire California data (as seen in Figure 5(b) that shows the distribution of land cover categories in California). In *DSCanada* we include the entire state of Yukon without any MODIS land cover mask as in this region almost all locations are forest, savanna or shrub. Also note that there are non-fire related changes in the forest, savanna and shrub MODIS categories such as logging and conversion to agriculture which are not covered in the fire polygons. Such changes will be incorrectly considered as false positives. However, we expect that this issue will impact performance of all algorithms similarly and will not change their relative performance.

4.2. **Evaluation Methodology.** The change detection algorithms assign a change score to each location, and the locations are ranked according to the descending order of their change score. The algorithm flags the top n ranked locations as change events and the lower ranked locations as unchanged. By computing the intersection with the validation data, we find the number of true positives (TP_n), false positives (FP_n), true negatives (TN_n) and false negatives (FN_n) as shown by Table 1. Our evaluation of the performance of the change detection algorithms is based on computation of the *precision* and *recall* that are two well-known metrics used to evaluate the performance of algorithms in information retrieval, machine learning and data mining [19] and given by:

$$\text{Precision, } p_n = \frac{TP_n}{TP_n + FP_n}$$

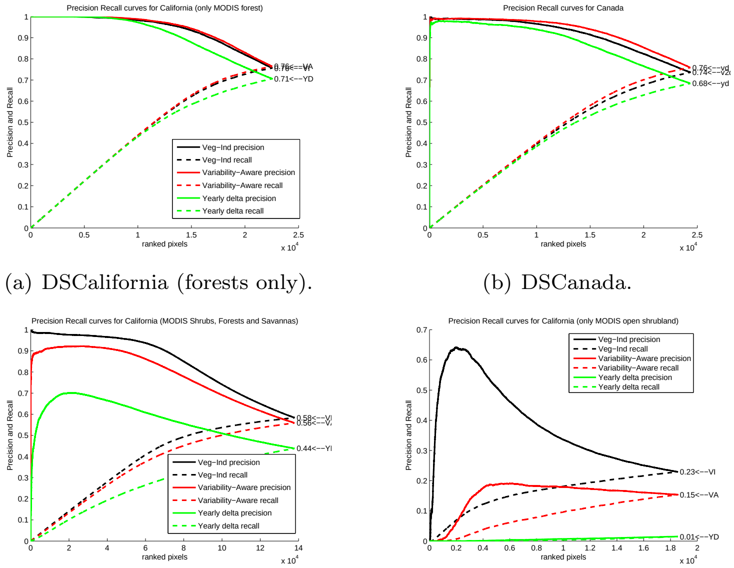

(a) DSCalifornia (forests only). (b) DSCanada.

(c) DSCalifornia. (d) DSCalifornia (open shrubs only).

FIGURE 6. Precision and Recall curves for three algorithms. *black* for VID, *red* for VD and *green* for YD

$$\text{Recall, } r_n = \frac{TP_n}{M}$$

To compare the relative performance of different techniques, we plot the precision and recall curve for the ranked list of pixels for the values $1 \leq n \leq M$. An ideal change identification algorithm should have a precision of 1 and a steadily rising recall from 0 to 1 as n increases from 1 to M. M in our case is the actual number of pixels inside fire polygons.

5. Discussion of Experimental Results

The performance of the three algorithms described in Section 3 is evaluated and analyzed in this section. The precision and recall plots for the data set on Yukon (Canada) and California (USA) are used to understand the accuracy differences between the three algorithms on different vegetations. *DSCalifornia* data set comprises of many different land cover types and to illustrate the effect of incorporating variability in the change detection algorithms, this data was subsetted based on MODIS landcover map and results are reported for only forests and shrubs in addition to the entire data for California. A comparison with Burned Area Product is also reported to further highlight the ability of the algorithms to identify forest fires from EVI data.

5.1. **Performance of proposed approaches.** Figure 6(a) and Figure 6(b) show the precision and recall curve for the three algorithms in California (only forests) and Yukon (Canada) respectively. The performance of the three algorithms is comparable in the data sets that comprise of only MODIS forest land cover category in California (USA) and in entire Yukon province (Canada). All three algorithms perform well and are able to identify fires in these area with high recall and precision, but the two algorithms with variability incorporated show slightly better results. The reason behind the good performance of all algorithms on the *DSCalifornia* data set is that forests have lower variability and typically have a stable EVI which has an abrupt decrease primarily in case of an actual land cover change. The decrease in precision occurs primarily because the algorithms identify other changes like logging that also show an abrupt decrease in EVI but as the ground truth is limited to forest fires these changes are considered as false positives. This limitation of the ground truth is further discussed in Section 6. The slight improvement in performance by modeling natural variability comes because of presence of some non-forested locations in the data set. We use the MODIS forest map in California for this data set but it is inaccurate and includes some shrubs and agriculture lands labeled as forests. The VD and the VID are more resilient to such misclassifications in land cover map. This is because farms are by nature highly variable due to shifts in cropping dates and other reasons, and as such get a high variability score and are therefore eliminated by algorithms that incorporate variability modeling, but are detected as changes in the YD because it gives these locations a higher change score. The locations in *DSCanada* data set have primarily MODIS forest, savanna and shrub category. The EVI signal for all locations in this data set has considerable homogenity and therefore VD and VID are only slightly better compared to YD. We see that VID performs slightly worse than VD on the *DSCanada*. This is because in the EVI time series in Canada, some observations have an unusually high EVI value. If this noise occurs in the first few years that are used to model the variability, the variability of that location will be high and as a result a change in the later years at that location will not get detected due to the high variability score. These outliers negatively impact results by incorrectly increasing variability for such locations and since VID uses variability modeling more strictly as compared to VD, it gets negatively impacted due to these outliers to a greater extent and shows a slightly worse performance than VD as seen in Figure 6(b).

The contrast between the performance of the algorithms becomes evident when we evaluate on the entire *DSCalifornia* data set. This is because this data set has multiple land cover types (including shrubs) and variability modeling becomes essential in the case of some of the land cover categories. Figure 6(c) shows the quantitative performance for the three algorithms on this data set. We notice that the YD algorithm performs significantly poorer than its counterparts that incorporate variability modeling. This is not surprising as shrubs form the dominant land cover type in the data set and they have high variability due to their higher sensitivity to climate variation (eg. precipitation). YD gives a high change score to many locations even if they are not burnt and thus leads to a poor precision on the composite data set (*DSCalifornia*). This fact is further illustrated in Figure 6(d), which shows the performance of the three algorithms on *DSCalifornia* with open shrubs only. The precision of YD on this data set is exceptionally poor and indicates that YD is not a good change detection algorithm for this land cover type though its performance is comparable to other algorithms in forests. Since the number of shrubs in the composite data set *DSCalifornia* is high, the poor precision of YD on this data set is explainable. The high variability of shrubs is also present in the first few years used to compute μ_{var} and thus the VD is able to perform better and shows a considerably improved performance over YD on *DSCalifornia*. Again, this fact is suported by the observation in Figure 6(d) where precision of VD is significantly better than that for YD. The VID algorithm that also models the spread of the variability distribution along with the μ_{var}, is able to work well across different land cover types. This is primarily because *DSCalifornia* (open shrubs only) has EVI time series with large variations in the spread of their variability scores (see Figure 3(b)). The VID score takes into account the information about the distribution's spread and therefore avoids several false alarms from this vegetation, while other algorithms make many mistakes on this particular vegetation. Low recall on the *DSCaliforina* (only open shrublands) data set is primarily due the fact that for many pixels inside fire polygons the EVI shows no significant change. Other bands of the spectrum have to be analyzed to be able to identify these burnt locations. An example EVI time series of such a location that was inside a fire polygon but has little change in EVI is shown in Figure 8. The red vertical line marks the date of fire. We see that the decrease in EVI was too small to be identified by these algorithms.

Another observation is that the VID algorithm has a much better performance in locations with low mean EVI (see Figure 7(a)). These vegetation types are typically stable and even at change, the decrease in EVI response in small in magnitude. Since VID models the standard deviation in the

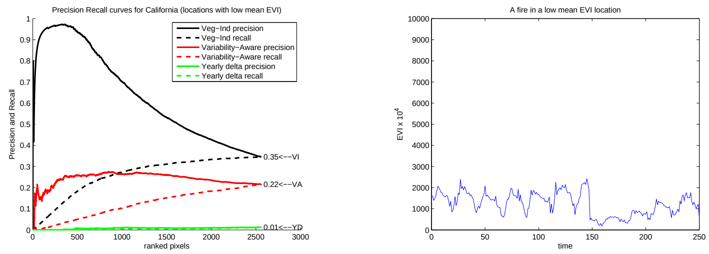

(a) Precision and Recall curves for three algorithms on low EVI (less than 1300). *black* for VID, *red* for VD and *green* for YD.

(b) A location in California with low and stable EVI and a fire with small decrease in EVI.

FIGURE 7. VID performs better on Low EVI locations and correctly detects changes in such locations.

variability, these vegetation types have a low σ_{var} and therefore even smaller changes get identified. This is the case in fires in the open shrub land cover for California which typically occur in the locations with extremely stable low mean EVI values (an example is shown in Figure 7(b)). It is observed that the VID score is independent of the the magnitude of the original time series and has therefore has comparable performance across vegetation types with different mean EVI.

5.2. **Comparison with Burned Area Product.** We use the output of the Burned Area Product to evaluate its performance on *DSCalifornia* and *DSCanada*. The differences in the performance of this product and our algorithms occur due to two reasons: (1) the change detection mechanism used in this product is different from our approaches and (2) the data set used for the generation of this product is thermal band instead of EVI. The Burned Area Product indicated 34,986 locations to have burned in 2006-2008 in *DSCalifornia*. Out of these 24,890 are inside the polygons. The precision and recall is 71.1% and 18.1% for Burned Area Product on *DSCalifornia*. In Canada, the performance of Burned Area Product is better than that in California. The Burned Area Product reports 15,005 pixels burnt in 2004-2008 out of which 13,513 are in polygons. The precision and recall is 90% and 55.5% respectively on *DSCanada*. For the same precision as Burned Area, the YD algorithm has a similar recall on both *DSCalifornia* and *DSCanada*. This indicates that YD has a comparable performance to the technique by Roy et al. [18] that is used to generate the Burned Area Product. For a similar precision on the two data sets as the Burned Area, recall for the VID algorithm is around 50% on the *DSCalifornia* and 60% on *DSCanada*. The results on *DSCalifornia* are significantly better for the VID algorithm over the YD and Burned Area. This is because this data set has multiple vegetation types present in it and VID that incorporates the natural variation of EVI time series in the change detection paradigm performs better. The same idea of incorporating variability in the change detection framework can potentially be used with the Roy et al. [18] approach and improve their change detection accuracy. Furthermore, we notice some complementarity between the change events detected by the two approaches. This is primarily due to the different data set (thermal band) used by Burned Area Product. Several changes like Figure 8 which are not prominent in EVI signal are detected in Burned Area Product.

6. Concluding Remarks

In this paper, we described two novel time series change detection algorithms that can be used to identify abrupt vegetation loss and extend the Yearly Delta algorithm by introducing the concept of natural variation in EVI response of a location. The results of the study demonstrate the importance of modeling the natural variation in the vegetation response for each location for accurately estimating the significance of the change in EVI signal. The evaluation of the proposed method is done

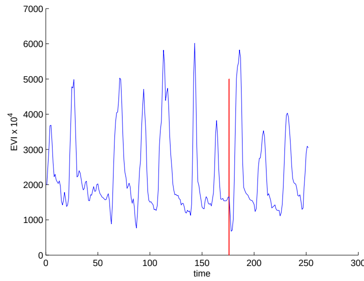

FIGURE 8. A burned location in California correctly identified by Burned Area Product and went undetected by our approaches because it shows little change in EVI signal. The vertical red line marks the change date.

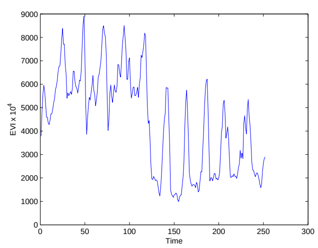

(a) Change of forest to other vegetation phenology

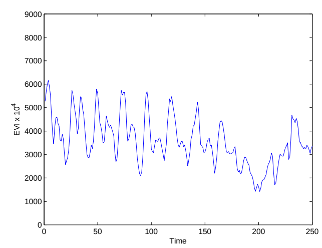

(b) Gradual decrease in time series

FIGURE 9. EVI time series for a pixel not present in fire polygons while the time series indicate change.

quantitatively using the validation data on forest fires from California state (USA) and Yukon state (Canada). The evaluation results demonstrate the ability of our proposed approach in identifying occurrence of forest fires from a remote sensing vegetation dataset (Enhanced Vegetation Index) with high accuracy. These algorithms are computationally fast (3000 timeseries get processed per second using a MATLAB implementation on a desktop), making it possible to process the entire globe at 1 km spatial resolution in less than a day on a desktop computer. Since computation for each time series is independent, the algorithms are easily parallelized on a current generation multi-core computer. In the following, we discuss the limitations of the current work and possible directions for future research.

- **Limitation of evaluation methodology:** A careful look at the false positives of the Vegetation-Independent Yearly Delta algorithm reveals that it finds many changes that do not correspond to fires. This is because an abrupt decrease in EVI is also caused by other forest disturbances such as logging, floods, conversion to agriculture, etc (Figure 9(a)). In addition, it often finds gradual decreases in EVI that might occur due to slow forest degradation such as in Figure 9(b). These locations, though genuine forest cover disturbances, are considered false positives because they are absent in the ground truth which is restricted to forest fires. Lack of exhaustive ground truth for land cover change is a serious problem with evaluation of forest monitoring algorithms. A fair evaluation is possible if an additional characterization step to identify the type of the change is included. This is particularly challenging because several types of changes often have similar EVI signature. For example, fires and deforestation often show the same characteristic abrupt decrease in EVI. The use of

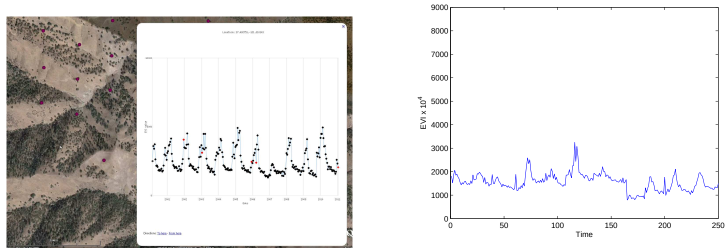

(a) EVI time series for a pixel in shrub location that is not present in fire polygons. The time series indicate an unusually low vegetation response for an year.

(b) EVI time series for a pixel that was inside a fire polygon and not detected by the VID algorithm

FIGURE 10. Example time series to illustrate the need for including climate variables like precipitation in change detection framework

complex spatial and temporal structures present in the EVI data in conjunction with other data sets (for eg. thermal band) could help distinguish between such changes.

- **Modeling of variations due to climate variables:** Figure 10(a) shows an example of a location that shows unusually low values for the vegetation index in one fo the years. This signature was present in many pixels in shrublands in California and appears to be result of a drought-like condition. These pixels are flagged as changes by our algorithms since they correspond to sudden drop in vegetation, but they are considered as false positives in the context of detecting fires or deforestation events. Figure 10(b) is an example of a pixel that was not identified as a change, though it was in the ground truth. This pixel had a high variability score perhaps because the vegetation is highly sensitive to precipitation etc., hence the relatively smaller decrease was missed. One approach to correctly handle such cases, is to model variations due to changes in climate and incorporate it in the change detection paradigm. Hammer et al. [8] use rainfall in the month as a dependent variable in the regression equation for monthly NDVI to account for changes in rainfall. The extra-seasonal relationship between rainfall and NDVI is captured by this term. In addition, a long term trend in amount of rainfall leading to a trend in NDVI can be captured.

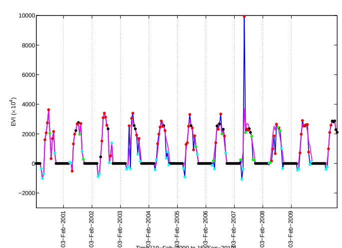

FIGURE 11. EVI time series for a pixel not present in fire polygons but was detected as change due to a noisy value in year 2007.

- **Noise Reduction:** Noise in the EVI time series poses a significant challenge to any change detection algorithm. For example, all algorithms in this study falsely identify changes in cases such as in Figure 11 where there is a noisy observation in year 2007 that increases the mean annual EVI for that year. There is a need to design a noise reduction technique that is cognizant of characteristics of remote sensing datasets and that utilizes the information about quality of observations, cloud and aerosol conditions that are available with the data. Remote Sensing community has developed many noise reduction techniques to reduce the impact of noise in these data sets [9]. However, since these techniques were not designed to account for a possibility of an abrupt change, these smoothing-based techniques tend to distort the actual change point. It is therefore not possible to use an off-the shelf noise reduction technique and new noise reduction techniques need to be developed that are suitable in the context of change detection.

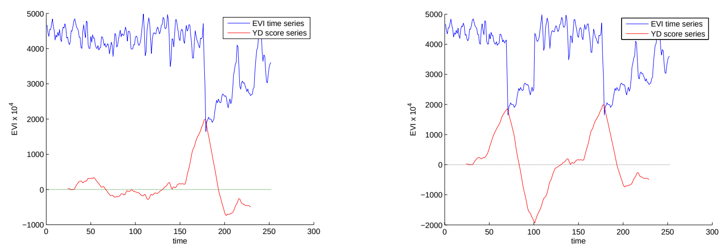

(a) EVI time series (in *blue*) for a location with a single fire and the corresponding YD score time series (in *red*).

(b) EVI time series (in *blue*) for a location with two fires and the corresponding YD score time series (in *red*).

FIGURE 12. EVI time series and corresponding YD time series for single and multiple changes

- **Identifying multiple changes in time series:** Another limitation of the proposed approach is that it can find only a single change in a time series. The ability to find multiple changes will become critical as these time series are increasing in time dimension with more satellite data being collected. The change detection framework needs to be extended to allow for finding multiple changes in the time series. Thus instead of assigning each location a single change score, time steps of a location should get flagged as changes if they correspond to local maximas that are higher than the score threshold. As an example, Figure 12(a) shows the EVI and YD score time series for a single change. The peak in the YD score time series corresponds to the change point. Figure 12(b) shows the EVI time series for a location with two fires. The YD score time series shows two peaks that are separated in time and peak at the time of fire. Both these peaks should be flagged as changes and the location will have two change events.
- **Model Selection:** The change detection algorithms described in this paper use the previous year as a model for EVI values of the current year. A more robust model is built using the mean or median of the EVI values in the previous k years. This is especially useful in eliminating false alarms due to noise in data or climate variations such as seen in Figure 4. In this figure the vegetation response is high for two years and the YD score will be high for the next year that has a low response. But if the median score for the previous 5 years was used the score will be low and not get falsely detected as change.

References

[1] Land Processes Distributed Active Archive Center. `http://lpdaac.usgs.gov`.

[2] L. O. Anderson, Y. Malhi, L. E. O. C. Aragão, R. Ladle, E. Arai, N. Barbier, and O. Phillips. Remote sensing detection of droughts in Amazonian forest canopies. *New Phytologist*, 187(3): 733–750, 2010. doi: 10.1111/j.1469-8137.2010.03355.x.

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

[4] 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.

[5] S. Boriah, V. Mithal, A. Garg, V. Kumar, M. Steinbach, C. Potter, and S. Klooster. A comparative study of algorithms for land cover change. In *Conference on Intelligent Data Understanding*, 2010.

[6] 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.

[7] M. A. Friedl, D. Sulla-Menashe, B. Tan, A. Schneider, N. Ramankutty, A. Sibley, and X. Huang. MODIS Collection 5 global land cover: Algorithm refinements and characterization of new datasets. *Remote Sensing of Environment*, 114(1):168–182, 2010.

[8] D. Hammer, R. Kraft, and D. Wheeler. Forma: Forest monitoring for actionrapid identification of pan-tropical deforestation using moderate-resolution remotely sensed data. *Working Papers*, 2009.

[9] J. Hird and G. McDermid. Noise reduction of ndvi time series: An empirical comparison of selected techniques. *Remote Sensing of Environment*, 113(1):248 – 258, 2009.

[10] R. Kennedy, Z. Yang, and W. Cohen. Detecting trends in forest disturbance and recovery using yearly landsat time series: 1. landtrendr–temporal segmentation algorithms. *Remote Sensing of Environment*, 2010.

[11] A. Koltunov, S. L. Ustin, G. P. Asner, and I. Fung. Selective logging changes forest phenology in the Brazilian Amazon: Evidence from MODIS image time series analysis. *Remote Sensing of Environment*, 113(11):2431–2440, 2009.

[12] J. Kucera, P. Barbosa, and P. Strobl. Cumulative sum charts-a novel technique for processing daily time series of modis data for burnt area mapping in portugal. In *Analysis of Multi-temporal Remote Sensing Images, 2007. MultiTemp 2007. International Workshop on the*, pages 1–6. IEEE.

[13] D. Lu, P. Mausel, E. Brondízio, and E. Moran. Change detection techniques. *International Journal of Remote Sensing*, 25(12):2365–2401, 2003.

[14] 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.

[15] 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 Transactions on Intelligent Systems and Technology*, (in press), 2011.

[16] R. Radke, S. Andra, O. Al-Kofahi, and B. Roysam. Image change detection algorithms: a systematic survey. *IEEE Transactions on Image Processing*, 14(3):294–307, 2005.

[17] N. Ramankutty, H. K. Gibbs, F. Achard, R. Defries, J. A. Foley, and R. A. Houghton. Challenges to estimating carbon emissions from tropical deforestation. *Global Change Biology*, 13:51–66, January 2007.

[18] D. P. Roy, P. E. Lewis, and C. O. Justice. Burned area mapping using multi-temporal moderate spatial resolution data–a bi-directional reflectance model-based expectation approach. *Remote Sensing of Environment*, 83(1-2):263–286, 2002.

[19] P.-N. Tan, M. Steinbach, and V. Kumar. *Introduction to Data Mining*. Addison-Wesley, Boston, MA, 2006.

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