

A New Data Mining Framework for Forest Fire Mapping

Xi C. Chen[†], Anuj Karpatne[†], Yashu Chamber[†],
 Varun Mithal, Michael Lau, Karsten Steinhäuser,
 Shyam Boriah, Michael Steinbach, Vipin Kumar
 Dept. of Computer Science & Engineering
 University of Minnesota
 {chen, anuj, chamber, mithal, mwla, ksteinha,
 sboriah, steinbac, kumar}@cs.umn.edu

Christopher S. Potter,
 Steven A. Klooster
 NASA Ames Research Center
 chris.potter@nasa.gov,
 sklooster@gaia.arc.nasa.gov

Teji Abraham,
 J.D. Stanley,
 Juan Carlos Castilla-Rubio
 Planetary Skin Institute
 {teabraha, jdstanle, juanccas}
 @cisco.com

Abstract—Forests are an important natural resource that support economic activity and play a significant role in regulating the climate and the carbon cycle, yet forest ecosystems are increasingly threatened by fires caused by a range of natural and anthropogenic factors. Mapping these fires, which can range in size from less than an acre to hundreds of thousands of acres, is an important task for supporting climate and carbon cycle studies as well as informing forest management. Currently, there are two primary approaches to fire mapping: field- and aerial-based surveys, which are costly and limited in their extent; and remote sensing-based approaches, which are more cost-effective but pose several interesting methodological and algorithmic challenges. In this paper, we introduce a new framework for mapping forest fires based on satellite observations. Specifically, we develop unsupervised spatio-temporal data mining methods for Moderate Resolution Imaging Spectroradiometer (MODIS) data to generate a history of forest fires. A systematic comparison with alternate approaches in two diverse geographic regions demonstrates that our algorithmic paradigm is able to overcome some of the limitations in both data and methods employed by prior efforts.

I. INTRODUCTION

Forest fires, which range in size from less than an acre to hundreds of thousands of acres, can be caused by both natural (e.g. lightning) or anthropogenic factors. Fires constitute a major component of terrestrial ecosystem disturbances every year, therefore accurate and low-cost fire mapping methods are important for understanding their frequency and distribution [26]. While monitoring fires in near-real time as they happen is critical for operational fire management, mapping historical fires in a spatially explicit fashion is also important for a number of reasons (recently highlighted in [16, 26]), including: (1) climate change studies – e.g., studying the relationship between rising temperatures and frequency of fires; (2) fuel load management – forest managers need a historical fire record when deciding where to conduct controlled burns; and (3) carbon cycle studies – quantifying how much CO₂ is emitted by fires is critical for emissions reduction efforts such as UN-REDD [32].

There are two primary approaches to mapping forest fires: field- and aerial-based studies, which allow detailed observa-

tion of land cover changes but are limited in spatial extent and temporal frequency [20] because of high cost; and remote sensing-based techniques, which offer the most cost-effective data for mapping fires because satellite observations – such as from NASA’s Moderate Resolution Imaging Spectroradiometer (MODIS) sensor – are obtained globally with regular, repeated coverage; here we focus on the latter. Such datasets have both temporal and spatial dimensions, and there are two main ways to address the problem.

On the one hand are approaches that focus on the temporal aspect, wherein fires are mapped based on time series analysis (e.g., [23, 28]). These types of methods usually take into consideration properties such as seasonality, variability and temporal coherence in a given time series. On the other hand are approaches that treat the data as a sequence of snapshots, and image processing-based methods (e.g., [7, 17]) are used to detect burned areas. Such methods mainly leverage the spatial properties inherent in the data, for instance, the fact that burned pixels tend to cluster together. Recently, techniques have been developed for land cover change detection that utilize both spatial and temporal properties [15, 19, 22] to take advantage of autocorrelation structures present along *both* dimensions. However, all of these are faced with a number of data, algorithmic, and computational challenges associated with analyzing remote sensing data including the presence of noise and outliers, incompleteness of signals, high natural variability and seasonality, influence of climatic factors, and availability of multiple spatial and/or temporal scales. In the case of fire mapping, additional factors include potential obstruction of the signal due to smoke and the similarity of the signal relative to other types of changes and events. Thus, while numerous efforts have mapped fires at regional and local scales [4, 6, 9, 25, 29, 30], only two spatially explicit wall-to-wall efforts exist that regularly map fires at global scale: the MODIS Active Fire (AF) [18] and Burned Area (BA) products [15].

AF is based solely on thermal anomalies and by itself is not sufficient to effectively map forest fires; it produces too many false positives because not all thermal anomalies are associated with fires while also missing a large portion of fires because

[†] These authors contributed equally to this work.


the signal in the composite data is not strong enough. BA is a more sophisticated method (see [15] and Section III-B) that partitions the task into groups based on a land cover map and identifies forest fires based on both temporal signals as well as spatial context. In particular, it uses both AF and a vegetation index computed from surface reflectance as input data to generate labeled training pixels within each land cover type; each class is further refined using information about nearby pixels in tropic and sub-tropic regions. Using these labeled samples as priors, a Bayesian model is then trained to distinguish burned from unburned pixels in a classification framework. BA generally provides much better performance than AF alone; however, it still has several notable limitations including sensitivity to false positive signals in AF and the reliance on a global land cover map, which is known to be inaccurate due to limitations of its own [10, 12].

In this paper, we introduce a new spatio-temporal data mining framework for forest fire mapping that is both robust and efficient. It shares some features with BA in that our proposed approach also exploits both temporal and spatial structure in the data and combines multiple sources of information – AF and a vegetation index – but it differs in several important respects: (1) it is *unsupervised* and therefore does not require labeled training data; (2) it does *not* rely on a global land cover map; and (3) it is *more robust* to noise and lower-quality signals in the input data. Using independently generated reference data for validation, we systematically evaluate our approach as well as alternate methods in two diverse geographic regions, and we examine how our approach is able to overcome some of the limitations of the BA algorithm.

The remainder of the paper is organized as follows: Section II briefly describes the input datasets. Section III discusses related work, including the BA algorithm, and in Section IV we introduce our new approach. Section V describes the experimental setup and validation data followed by results and discussion in Section VI. Section VII provides concluding remarks and directions for future work.

II. DATA

Global remote sensing datasets are available from a variety of sources at different resolutions. The proposed fire mapping framework is based on two remotely sensed composite data products from the MODIS instrument aboard NASA’s Terra satellite, which are available for public download [33]. Specifically, we use the Enhanced Vegetation Index (EVI) from the MODIS 16-day Level 3 1km Vegetation Indices (MOD13A2) and the Active Fire (AF) from the MODIS 8-day Level 3 1km Thermal Anomalies & Fire products (MOD14A2). EVI essentially measures “greenness” (area-averaged canopy photosynthetic capacity) as a proxy for the amount of vegetated biomass at a particular location (see Figure 1(a) for an example). AF is a basic fire product designed to identify thermal anomalies from the middle infrared spectral reflectance bands [18] and is used heavily in operational situations by fire-fighting agencies around the world. In order to separate forests from other land cover types, we use the MODIS Vegetation



(a) A sample EVI time series from the (b) A photograph of the fire in progress. Source: Associated Press

Fig. 1: The Basin Complex fire, which was started by lightning near Big Sur, CA in June 2008, consumed more than 160,000 acres before it merged with another fire and over \$120M was spent fighting it.

Continuous Fields (VCF) dataset (MOD44B), which provides the percent tree cover for every pixel. MODIS Level 3 products are provided on a global 1km sinusoidal grid in $10^\circ \times 10^\circ$ tiles. For this study, we focus on subsets of the data corresponding to geographical regions based on the available validation data (see Section V).

III. RELATED WORK

Fire-related data products broadly fall into two categories: *active fire* products, which capture the location and intensity of fires burning at the time of observation (the prototypical example being AF, see Section II); and *burned area* products, which map areas that were burned by fires based on historical observations. Here we discuss two approaches for mapping burned areas, which require more sophisticated analysis methods and are the topic of this paper.


A. The V2DELTA Algorithm

Mithal et al. [24] presented a time series change detection algorithm that incorporates *natural variation* into the change detection framework. The algorithm, called V2DELTA, identifies abrupt forest disturbances using EVI as an input. More specifically, the V2DELTA algorithm compares a drop in EVI with the variability in a fixed “training” window, thereby providing a mechanism to ascribe significance to any given drop. This relies on the assumption that EVI values in the initial window were not affected by a land cover change, thus enabling the algorithm to differentiate abrupt changes from naturally occurring vegetation changes.

While V2DELTA identifies a broad class of disturbances [24], it is not designed to distinguish fires from other land cover changes (e.g., droughts). Figure 2(a) shows a sample time series which was incorrectly identified as burned – the loss in EVI was actually due to logging. Moreover, the fire mapping task poses specific challenges that affect the algorithm’s efficacy, especially for time series that recover quickly (within a few months). Figure 2(b) shows one such example.

B. The BA Algorithm

The burned area approach (henceforth called BA) recently presented by Giglio et al. [15] is a state-of-the-art methodology in the earth science research community for identifying regions



(a) A pixel in California incorrectly identified as burned by V2DELTA. (b) A 2005 California fire event not detected as burned by V2DELTA because of fast recovery in “greenness”

Fig. 2: Illustrative examples of limitations in V2DELTA.

burned by fire. The overall approach can be viewed as a semi-supervised Bayesian classification method with two classes: *burned* and *unburned*. The technique builds on key concepts and ideas developed over several years by Giglio et al. [14] and others [11, 13, 21, 27]. The key steps of the algorithm are outlined below:

- 1) Representative sets of samples for the *burned* and *unburned* classes are constructed. The sample pixels for each class are discovered using conservative heuristics which label pixels as unburned or burned if they pass a set of conditions.
- 2) The *burned* class is further enriched with closely related pixels from the dataset, while the *unburned* class is refined by pruning pixels that are geographically close to *burned* training pixels.
- 3) A statistic that estimates the daily loss in vegetation (ΔVI) is computed for all training pixels.
- 4) The conditional probability distribution of the vegetation loss statistic is estimated for both the *burned* as well as the *unburned* class, i.e., $P(\Delta VI|burned)$ and $P(\Delta VI|unburned)$.
- 5) Bayes’ Rule is applied to obtain the posterior probability of a pixel belonging to the *burned* class.

The BA algorithm is run on a regular basis using the latest spectral reflectance and AF inputs. The output is released by the University of Maryland as a product called MODIS Direct Broadcast Monthly Burned Area Product (MCD64A1). Two versions of the BA product are available: one that only uses input data deemed to be of high quality and one that also incorporates lower-quality input data (which we call BAHighQ and BALowQ, respectively). Although the former is the one that is widely used, our experimental results show that the latter based on lower-quality inputs can produce better results (see Section VI).

IV. PROPOSED APPROACH

Forest fires lead to the burning of vegetation (trees and shrubs) and the emission of large amounts of thermal energy close to the land surface. Thus, a forest fire typically exhibits simultaneous changes (at a given location and time) in both the greenness as well as the thermal anomaly signals. We utilize these properties to identify fire events by performing

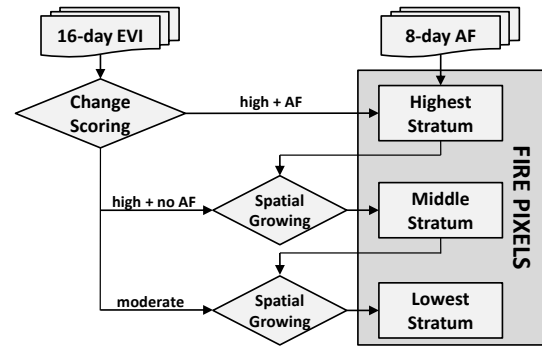


Fig. 3: Flowchart illustrating the proposed framework for mapping forest fires.

an integrated analysis in both the EVI and AF datasets. In addition, fires exhibit particular spatio-temporal properties which can be harnessed to further improve fire mapping. The full workflow consists of generating stratified sets of pixels based on confidence of change exhibited in the EVI and AF signals.

In particular, we present a framework (Figure 3) for mapping forest fires wherein the multiple strata signify varying degrees of observability in the available datasets and confidence in detecting them. To obtain the highest stratum of detected fire events, we employ multiple complementary scoring mechanisms using both EVI and AF data. This stratum is expanded by including very similar events in close proximity to form the middle stratum. The lowest stratum is generated by including loosely similar events in a spatial window around the other two strata.

A. EVI Scoring Mechanisms

Forest fires are often characterized by a sudden decrease in the EVI time series. Sometimes, these drops will persist for a few years (e.g., fires in boreal evergreen forest) and sometimes they only last for several months (e.g., fires in tropical rainforests). Generally, they are significant in both absolute value and relative value compared with the normal variations attributed to climatic seasonality and sensor noise. Below, we define several scoring mechanisms, each of which captures some aspects of change and variation for a given time series.

a) *K-month Delta (KD)*: KD is designed to score the changes which persist for a long time (k months). It accounts for the natural variability present in the vegetation, which is specific to a particular region. By modeling EVI as the combination of yearly trend and Gaussian noise, we assume that inter-annual variation (IAV, defined below) follows a Gaussian distribution:

$$IAV(t) = \mu(EVI(t - sl, K)) - \mu(EVI(t, K))$$

where sl is the number of time steps in one annual EVI segment (23 in our case) and K is the window size of segments being compared (12 months).

Therefore, the KD score at time step t is given by

$$KD(t) = \frac{IAV(t)}{\sigma}$$

which is the z-score of IAV. Here, σ is estimated based on the data in a four year window preceding t using bootstrapping (this makes the algorithm more robust than V2DELTA).

b) Local Instant Drop (LID): LID scores the instantaneous drop in EVI to identify fires that recover too quickly to be captured by the KD score. The algorithm accounts for the seasonal context in EVI to improve the robustness of the scoring algorithm. Specifically, the LID score is given by

$$LID(t) = \frac{EVI(t-1) - EVI(t+1)}{NVar}$$

where $NVar$ is the largest drop that occurs in the temporal neighborhood (of size 3) in the previous two year history. In order to account for the seasonality of EVI, only the time steps within a small window (of size 1) around a given time step in the previous years are considered.

LID is not as stable as KD and therefore generally has a higher threshold to guarantee the detection accuracy. Additionally, it is a good filter to remove false positives detected by KD due to other types of vegetation changes, such as drought.

c) Near Drop (ND): ND is designed to enforce a minimum absolute change in EVI when a fire happens. As the only score which reflects the real amount of drop in EVI, ND is well-suited as a filter in our framework. In particular, ND is calculated as

$$ND(t) = \mu(EVI(t-k, K)) - \mu(EVI(t+1, K))$$

where $K = 3$ steps.

d) Integrating Multiple Scoring Mechanisms: It is evident from the discussion above that the scoring mechanisms serve complementary purposes and capture different characteristics of the fire event, which are distinct in nature. Hence, they offer a possibility for developing a fire detection framework which utilizes the orthogonal aspects of each scoring mechanism in an integrative manner.

B. Initial Pixels

A forest fire exhibits large instantaneous drop; thus, a high ND score (≥ 500) is a filtering criterion that must be met by all pixels. A large KD score (≥ 3) is representative of a forest fire event when coupled in conjunction with a moderate LID score (≥ 1), which helps in rejecting other land cover changes that show high KD score but are not associated with fires (low LID score). It can be used to detect fire which has long-term effect on the forests. A large LID score (≥ 4) is in itself a good indicator of forest fire events. It is useful in detecting the fires that occur in the regions where greenness recovers quickly. Events which satisfy either of these two scoring criteria are considered as initial pixels (**highest stratum**) and are further used in subsequent steps of our fire mapping framework.

C. Spatial Growing

The AF signal often fails to detect forest fire events which do not register a thermal anomaly because of smoke or satellite overpass timing. Thus, the initial pixels might suffer from low coverage. To overcome this challenge, we exploit the inherent spatio-temporal autocorrelation of forest fire events to increase coverage. Since events corresponding to the same forest fire occur in close proximity of space and time, we exploit this property by searching for fire events around the initial pixels classified as forest fires by the scoring mechanism above. In the current framework, we consider the 24 spatial neighbors in a 5×5 spatial grid around the initial pixels, with a temporal constraint of being within one time step from the change time of the initial fire event. We then apply our scoring mechanism on the new pool of candidate events with exactly the same scoring criteria as we used for detecting initial forest fire events. We iteratively grow in a spatial neighborhood to exhaustively detect forest fire events. They represent forest fire events (**middle stratum**) which have sharp fire characteristics in the EVI signal but were not initial pixels because of the absence of AF signal.

Because of the presence of noise in EVI as well as cases where EVI loss is small, there are a number of forest fire events which do not exhibit strong characteristics of fire in our scoring mechanisms and thus go undetected. Simply lowering the threshold in initial pixel detection will decrease the robustness of our approach on the noise of AF. Therefore, we exploit the spatial autocorrelation of forest fires to discover such events.

In particular, we exploit this property to create another level of forest fire events (**lowest stratum**) with a relaxed scoring criteria indicating a lower confidence. We accept events to be part of the lowest stratum if they exhibit a positive ND score and either a moderately large LID score (≥ 2), or a moderately large KD score (≥ 2.5) in conjunction with a moderate LID score (≥ 0.8). Thus, we iteratively grow in a spatial neighborhood (5×5 grid) to exhaustively include any probable fire events. This improves our coverage of forest fires but is likely to include some false positives.

V. EVALUATION SETUP

We examine the performance of AF, BA and the proposed approach in two ecologically varied regions, namely, the US states of California and Georgia. These geographic areas represent diverse regions with widely differing variability, land cover types, geography and noise characteristics. The following describes the validation data used in this study and provides a brief overview of the evaluation methodology.

A. Validation Data

For each region considered in our evaluation, we obtained fire validation data from government agencies responsible for monitoring and managing forests and wildfires. The validation data is in the form of fire perimeter polygons, each of which is associated with the time of burning. Table I summarizes the regions studied in this paper and the respective sources of

Region	References	Positives	Negatives
California (US)	[1, 8]	4597	597208
Georgia (US)	[2, 3]	2003	425528

TABLE I: Regions studied in this paper and their respective sources of historical fire validation data.

Validation Data	Predicted		
		Fire	No Fire
	Fire	<i>TP</i>	<i>FN</i>
No Fire	<i>FP</i>	<i>TN</i>	

TABLE II: Confusion matrix.

validation data. Although government agencies make their best effort in documenting historical fires, fire perimeter datasets are neither complete nor without error due to the finite resources available to any agency. However, inaccuracies and incompleteness are represented only in a small portion of the validation data, and these datasets are still useful for quantitatively comparing methods across large spatial regions. The AF, BA and EVI datasets are georeferenced by the latitude and longitude values of their pixel centers. We consider an event to be positive if the corresponding pixel lies inside a polygon. Similarly, an event is considered to be unburned (forming the negative class) only if the entire pixel is outside a polygon. The remaining pixels (which partially overlap polygon boundaries) are discarded from the evaluation framework to avoid ambiguity.

Since our primary focus is on detecting forest fires, we utilize the MODIS Vegetation Continuous Fields (VCF) dataset (MOD44B) which contains the percentage tree cover information. We only consider pixels with high percentage tree cover (≥ 20) in our evaluation scheme, a procedure commonly used in the earth science domain to separate forests from non-forests [5, 15].

B. Evaluation Methodology

In this paper, we use *precision* and *recall* as evaluation metrics for quantitatively comparing the performance of AF, BA and the proposed approach. These two well-known metrics are used to evaluate the performance of algorithms in information retrieval, machine learning and data mining [31]. Each algorithm provides a set of positive and negative events that it detects, which is validated using fire perimeter polygons to obtain the number of true positives (*TP*), false positives (*FP*), false negatives (*FN*) and true negatives (*TN*) for each algorithm, as shown in Table II. Note that a *TP* event means that the pixel lies inside a polygon and the time of change agrees with the polygon. The precision and recall values for each algorithm are then given by:

$$\text{Precision, } p = \frac{TP}{TP + FP} \quad \text{Recall, } r = \frac{TP}{TP + FN}$$

VI. EXPERIMENTAL RESULTS & DISCUSSION

Tables III and IV contain validation results of our proposed fire mapping framework along with other algorithms discussed in this paper for the two different geographic regions. In the following paragraphs we discuss their performance as well as relative strengths and weaknesses in detail, beginning with some general observations.

Algo	Precision	Recall	TP	FP	FN
Highest Stratum	0.994	0.806	3748	21	901
Middle Stratum	0.990	0.844	3922	38	727
Lowest Stratum	0.944	0.906	4210	249	439
BAHighQ	0.990	0.926	4307	44	342
BALowQ	0.988	0.926	4307	52	342
AF	0.624	0.913	4246	2563	403
V2DELTA	0.054	0.960	4462	77827	187

TABLE III: Evaluation results for California.

Algo	Precision	Recall	TP	FP	FN
Highest Stratum	0.985	0.475	956	15	1056
Middle Stratum	0.986	0.601	1209	17	803
Lowest Stratum	0.973	0.707	1422	39	590
BAHighQ	0.470	0.193	389	438	1624
BALowQ	0.700	0.674	1356	581	656
AF	0.147	0.644	1295	7522	717
V2DELTA	0.027	0.869	1749	64237	263

TABLE IV: Evaluation results for Georgia.

First, we note that our approach is more stable in the two geographic areas even though in California, BA performs slightly better. Since our approach takes into account the natural variation and spatial coherence, our framework is more reliable when the signal quality is not good. But when the signals are clean, the heuristic method designed by domain experts performs (slightly) better than ours.

Second, the quality flag in BA is not trustworthy since it represents the quality of input data not the output result. Unlike BA, the quality in the proposed approach reflects the strength of “fire” pattern in a time series and always guarantee a higher recall but lower precision for a lower stratum.

Finally, by comparing the third and seventh row in Tables III and IV, we find that our framework consistently outperforms the V2DELTA algorithm [24]. Although V2DELTA achieves a comparable recall in some cases, it tends to generate a large number of false positives which is not surprising given that it is a more general change detection method. In fact, it was this shortcoming that prompted the development of the bootstrapping procedure for more robust modeling of the historical EVI patterns and incorporation of the AF signal as an additional source of information. This point is illustrated in Figure 2(a), which shows the time series for a pixel that was incorrectly identified by V2DELTA as burned. However, with a combination of robust time series modeling and AF signal,

Data-set	Precision (3σ)	Recall (3σ)
AF + Noise(N)	0.9936 ± 0.0016	0.8166 ± 0.0006
	0.9936 ± 0.0015	0.8754 ± 0.0003
	0.7479 ± 0.0786	0.9571 ± 0.0011
AF + Noise(5N)	0.9933 ± 0.0015	0.8165 ± 0.0008
	0.9934 ± 0.0016	0.8754 ± 0.0008
	0.7558 ± 0.0631	0.9571 ± 0.0011
AF + Noise(10N)	0.9935 ± 0.0015	0.8165 ± 0.0006
	0.9935 ± 0.0017	0.8753 ± 0.0004
	0.7411 ± 0.0849	0.9570 ± 0.0004

TABLE V: Performance of the proposed framework in California given “noise” in the form of additional (false positive) AF signals. Each group of three rows corresponds to the top three rows in Table III.

our new framework is able to correctly dismiss such pixels as unburned.

A. Comparison with BA


Looking at the rows three and four in Tables III and IV, we note that our proposed framework is more stable in different geographic areas than BA. From our results, we further note that the overall performance of BA is tied closely to that of AF. For example, in California (Tables III) AF performs quite well, and so does BA. However, in Georgia (Table IV) AF incurs 7522 false positives resulting in very poor precision, which in turn severely limits the performance of BA.

This phenomenon is partly because BA depends on AF to generate the training data. Although several heuristic steps have been used in BA to improve the accuracy of the training data, the quality of training data cannot be guaranteed when the precision of AF is too low (around 0.15 in Georgia). Compared to BA, the proposed framework is robust to the performance of AF because it does not simply search for sudden drops but focuses on detecting the patterns which represent “fire” in EVI time series.

In order to further illustrate the robustness of our approach to noise in AF, we performed additional set of experiments in California. We randomly inserted varying levels of false positives into the AF signal. The number of instances being added ranges from one to ten times the number of actual AF instances. Each experiment was repeated 30 times. The results, shown in Table V, are nearly identical to the main result (Table III) and suggest that our framework is extremely robust to the presence of false positive AF signals, unlike BA as seen in Table IV.

B. The Quality Flag compared with the ones used in BA

Although “high-quality” BA product is most frequently used in practice, in the results for Georgia (Table IV) we observe a peculiar case where both precision and recall are *higher* for the “low-quality” BA product. Unlike BA, our approach guarantees that lower level results have lower precision and higher recall. The main difference behind this observation is that the definition of “quality” in the two algorithms is not identical.



(a) TP in BA highQ

(b) FN in BA highQ but TP in BA lowQ

Fig. 4: A comparison of histograms for pixel reliability of BA suggests that the algorithm is sensitive to data quality.

When detecting “high-quality” results, BA uses only the high quality observations (input data) within a limited temporal window surrounding the fire date. Since values in the immediate (temporal) neighborhood of a fire tend to be of lower quality (due to obstruction from smoke), in some regions the algorithm does not have sufficient information and is thus unable to correctly classify such pixels. To test this hypothesis, we generated two histograms (Figures 4(a) and 4(b)) of the pixel reliability extracted from the Quality Assurance (QA) fields: the first is for the true positives of the high-quality product (left panel), and the second for the false negatives of the high-quality product which are true positives in the low-quality version (right panel). We find that the BA highQ events that were correctly identified were mostly high quality, while BA lowQ was able to take advantage of the additional information from lower-quality inputs and thus find a significantly larger number of burned pixels.

C. Limitations of the Proposed Approach

The forest fire mapping framework proposed in this paper faces limitations in a number of scenarios, leading to both false positives and false negatives. These include situations where (1) the vegetation rapidly recovers after a fire or there are multiple fires in short succession (FN), (2) the loss in vegetation is insignificant (FN), and (3) the vegetation has high natural variability (FP and FN). Each of these scenarios poses distinct challenges for our current fire detection framework. Additionally, if a fire polygon does not contain any pixel with an AF signal, it will not be detected by our framework (FN). However, we observed that such instances happen only in small polygons, and hence its effect on the performance of the proposed framework is insignificant.

We briefly present a discussion of the false positives and false negatives of the proposed approach. We had observed that a number of pixels exhibiting strong fire characteristics were missed by the fire perimeter polygons in California (Figure 5(a)). Such pixels were incorrectly counted as false positives in our evaluation scheme affecting the performance. In addition, we encountered a small number of false positives which exhibited other types of land cover changes (Figure 5(b)), and were incorrectly included in the highest stratum of detected points.

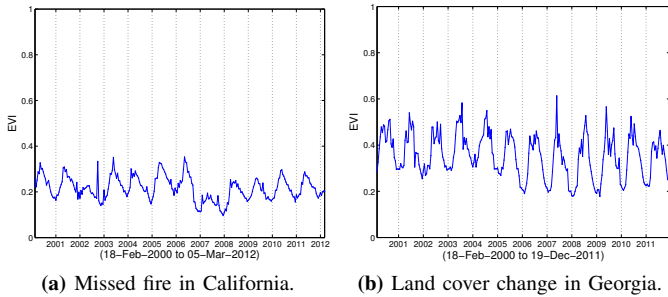




Fig. 5: False positives of the proposed approach.



(a) Vegetation with high variability in California (fire occurred in 2004).



(b) Weak fire characteristics in California (fire occurred in 2008). (c) Weak fire characteristics in Georgia (fire occurred in 2007).

Fig. 6: False negatives of the proposed approach.

Most false negatives of the proposed approach consist of vegetation types with high natural variability and noisy EVI data, which got poorly scored by the proposed algorithms and went undetected (Figure 6(a)). Other false negatives can be attributed to weak characteristics of fire in EVI signal, e.g., in California (Figure 6(b)) and Georgia (Figure 6(c)).

We are in effect considering *less information* than BA, yet we are able to achieve comparable performance. Thus, there is reason to believe that some of the limitations above will be addressed by increasing the spatial and temporal resolution to 250m and daily, respectively. In particular, we expect that smaller fires (250m data has sixteen times more detail than 1km data) and pixels that exhibit rapid recovery (because of compositing, neighboring time steps can be up to a month apart in the current 16-day data) can be detected with higher resolution data. We have not used higher resolution data in this paper since these are not standard MODIS products and hence require extensive processing to generate.

VII. CONCLUSION & FUTURE WORK

In this paper, we proposed a data mining framework for forest fire mapping. The proposed technique is unsupervised in nature and has the potential to be used globally, providing spatially explicit wall-to-wall coverage. We quantitatively showed that the technique performs comparably to the well-known BA algorithm in the state of California (US) and much better in Georgia (US); there are also complementarities between the two frameworks. We also showed that the proposed framework is highly robust to noise in one of its primary inputs, AF, which is known to have low precision.

While the current framework already performs relatively well in a variety of geographies, there are a number of interesting directions for future work. The data inputs used in this paper, EVI and AF, have temporal resolutions of 16 and 8 days, respectively. Both of these inputs can be computed on a daily basis, which is likely to provide much more information in many cases. This information can be exploited to identify the precise day of the fire and to ensure temporal coherence between neighboring pixels. Challenges with daily data include increase in the noise level and additional effort required to generate a daily EVI product (since this is not a standard MODIS product). Another potential extension is to use BA (and similar products) as an *input*, taking advantage of complementarities that exist between the frameworks.

ACKNOWLEDGMENT

This research was supported in part by the National Science Foundation under Grants IIS-1029711 and IIS-0905581, as well as the Planetary Skin Institute. Access to computing facilities was provided by the University of Minnesota Supercomputing Institute.

REFERENCES

- [1] California Department of Forestry and Fire Protection, Fire and Resource Assessment Program. <http://frap.fire.ca.gov>.
- [2] USDA Forest Service and US Geological Survey, Monitoring Trends in Burn Severity Project, Burned Area Boundaries Dataset. <http://mtbs.gov>.
- [3] National Fire Decision Support Center, Wildland Fire Decision Support System. National Fire Perimeters. <http://wfdss.usgs.gov>.
- [4] K. V. S. Badarinath, A. R. Sharma, and S. K. Kharol. Forest fire monitoring and burnt area mapping using satellite data: a study over the forest region of Kerala state, India. *International Journal of Remote Sensing*, 32(1):85–102, 2011.
- [5] S. Bandyopadhyay, P. Shyamsundar, and A. Baccini. Forests, biomass use and poverty in Malawi. *Ecological Economics*, 70(12):2461–2471, 2011.
- [6] Y. Bergeron, S. Gauthier, M. Flannigan, et al. Fire regimes at the transition between mixedwood and coniferous boreal forest in northwestern Quebec. *Ecology*, 85(7):1916–1932, 2004.
- [7] C. K. Brewer, J. C. Winne, R. L. Redmond, et al. Classifying and mapping wildfire severity: A comparison of methods. *Photogrammetric Engineering & Remote Sensing*, 71(11):1311–1320, 2005.
- [8] California Department of Forestry and Fire Protection. Fire management for California ecosystems, 1995.
- [9] M. Cochrane. *Tropical fire ecology: climate change, land use and ecosystem dynamics*. Springer, 2009.

- [10] G. M. Foody. Status of land cover classification assessment. *Remote Sensing of Environment*, 80:185–201, 2002.
- [11] R. H. Fraser, Z. Li, and J. Cihlar. Hotspot and NDVI Differencing Synergy (HANDS): a new technique for burned area mapping over boreal forest. *Remote Sensing of Environment*, 74(3):362–376, 2000.
- [12] M. A. Friedl, D. Sulla-Menashe, B. Tan, et al. MODIS Collection 5 global land cover: Algorithm refinements and characterization of new datasets. *Remote Sensing of Environment*, 114:168–182, 2010.
- [13] C. George, C. Rowland, F. Gerard, et al. Retrospective mapping of burnt areas in central Siberia using a modification of the normalised difference water index. *Remote Sensing of Environment*, 104(3):346–359, 2006.
- [14] L. Giglio, G. R. van der Werf, J. T. Randerson, et al. Global estimation of burned area using MODIS active fire observations. *Atmospheric Chemistry and Physics*, 6(4):957–974, 2006.
- [15] L. Giglio, T. Loboda, D. Roy, et al. An active-fire based burned area mapping algorithm for the MODIS sensor. *Remote Sensing of Environment*, 113(2):408–420, 2009.
- [16] J. Gillis. With deaths of forests, a loss of key climate protectors. *The New York Times*, October 1 2011.
- [17] I. Z. Gitas, G. H. Mitri, and G. Ventura. Object-based image classification for burned area mapping of Creus Cape, Spain, using NOAA-AVHRR imagery. *Remote Sensing of Environment*, 92(3):409–413, 2004.
- [18] C. O. Justice, L. Giglio, D. Roy, et al. MODIS-derived global fire products. In B. Ramachandran, C. O. Justice, and M. J. Abrams, editors, *Land Remote Sensing and Global Environmental Change*, pages 661–679. Springer, 2011.
- [19] S. Lhermitte, J. Verbesselt, I. Jonckheere, et al. Hierarchical image segmentation based on similarity of NDVI time series. *Remote Sensing of Environment*, 112(2):506–521, 2008.
- [20] D. Liu and S. Cai. A spatial-temporal modeling approach to reconstructing land-cover change trajectories from multi-temporal satellite imagery. *Annals of the Association of American Geographers*, 2011.
- [21] T. Loboda, K. O’Neal, and I. Csiszar. Regionally adaptable dNBR-based algorithm for burned area mapping from MODIS data. *Remote Sensing of Environment*, 109(4):429–442, 2007.
- [22] R. S. Lunetta, J. F. Knight, J. Ediriwickrema, et al. Land-cover change detection using multi-temporal MODIS NDVI data. *Remote Sensing of Environment*, 105(2):142–154, 2006.
- [23] V. Mithal, A. Garg, S. Boriah, et al. Monitoring global forest cover using data mining. *ACM Transactions on Intelligent Systems and Technology*, 2:36:1–36:24, 2011.
- [24] V. Mithal, A. Garg, I. Brugere, et al. Incorporating natural variation into time series-based land cover change detection. In *Proceedings of the 2011 NASA Conference on Intelligent Data Understanding (CIDU)*, pages 45–59, 2011.
- [25] M. Niklasson and A. Granström. Numbers and sizes of fires: Long-term spatially explicit fire history in a Swedish boreal landscape. *Ecology*, 81(6):1484–1499, 2000.
- [26] Y. Pan, R. A. Birdsey, J. Fang, et al. A large and persistent carbon sink in the world’s forests. *Science*, 333(6045):988–993, 2011.
- [27] D. P. Roy, L. Giglio, J. D. Kendall, et al. Multi-temporal active-fire based burn scar detection algorithm. *International Journal of Remote Sensing*, 20(5):1031–1038, 1999.
- [28] 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.
- [29] R. Somashekar, P. Ravikumar, C. Mohan Kumar, et al. Burnt area mapping of Bandipur National Park, India using IRS 1C/1D LISS III data. *Journal of the Indian Society of Remote Sensing*, 37:37–50, 2009.
- [30] B. Talon, S. Payette, L. Filion, et al. Reconstruction of the long-term fire history of an old-growth deciduous forest in Southern Québec, Canada, from charred wood in mineral soils. *Quaternary Research*, 64(1):36–43, 2005.
- [31] P.-N. Tan, M. Steinbach, and V. Kumar. *Introduction to Data Mining*. Addison-Wesley, Boston, MA, 2006.
- [32] United Nations. Collaborative Programme on Reducing Emissions from Deforestation and Forest Degradation in Developing Countries. <http://www.un-redd.org/>.
- [33] US Geological Survey and NASA. Land Processes Distributed Active Archive Center (LP DAAC). <https://lpdaac.usgs.gov/>.