The development and first validation of the GOES Early Fire Detection (GOES-EFD) algorithm

Alexander Koltunov^{1*},

Susan L. Ustin¹, Brad Quayle², Brian Schwind²,

Vincent G. Ambrosia³, Wei Li^{1,4}

¹Center for Spatial Technologies and Remote Sensing, University of California, Davis,
Veihmeyer Hall, One Shields Avenue, Davis, CA 95616, U.S.A.

²USDA Forest Service, Remote Sensing Applications Center (RSAC),
2222 West 2300 South, Salt Lake City, UT 84119 U.S.A

³California State University – Monterey Bay, NASA Ames Research Center,
Moffett Field, CA 94035, U.S.A

⁴College of Information Science and Technology,
Beijing University of Chemical Technology,
Beijing 100029, China

*To whom correspondence should be addressed:

E-mail: akoltunov@ucdavis.edu

KEYWORDS: geostationary; satellite; fire detection; wildfire detection; early wildfire detection; wildfire; detection timeliness; fire monitoring; remote sensing; GOES Early Fire Detection; GOES-EFD

Abstract

Decades of successful active fire mapping from space, have led to global informational products of growing importance to scientific community and operational agencies. In contrast, detecting fires from space faster than current conventional capabilities in the continental U.S. has not been considered attainable, except in remote, sparsely populated areas. We present a research prototype version of the GOES Early Fire Detection (GOES-EFD) algorithm focused on minimizing the time to first detection of a wildfire incident. The algorithm is designed for regional-scale surveillance and combines multitemporal anomaly tests developed in our previous work, contextual hot-spot tests, and dynamic event classification and tracking. The GOES-EFD version 0.4 was initially tested with 40-day summer 2006 data over central California. The algorithm identified most of large (final size >2 ha) wildfires within 30 min and 31% of the wildfires were detected before they were reported by the public. Under identical operation conditions, GOES-EFD 0.4 provided quicker initial detection than the temporally filtered operational WF-ABBA algorithm (version 6.1) and committed fewer false alarms. There is a substantial potential for further reducing detection latency and increasing reliability. Following the ongoing optimizations, tests, and integration in collaboration with the fire

19 management agencies and first responders, GOES-EFD could be deployed for
 20 regional scale real-time surveillance to complement existing fire identification
 21 methods.

22 List of acronyms

ABI	[GOES-R] Advanced Baseline Imager
BT	brightness temperature
BT ₄	brightness temperature in GOES band 2 (~4µm)
BT ₁₁	brightness temperature in GOES band 4 (~11µm)
c.c.	connected component
DDM	Dynamic Detection Model
EFD	early fire detection
GOES	Geostationary Operational Environmental Satellites
GOES-EFD	GOES Early Fire Detection
GVAR	GOES VARiable [format]
IADC	Iterative Anomaly Detection and Classification
INR	[GOES] Image Navigation and Registration
MODIS	Moderate Resolution Imaging Spectroradiometer
OCM	Operational Cloud Masking
RCD	Retrospective Cloud Detection
SCD	Single-Frame Cloud Detection
TIR	thermal infrared
WF-ABBA	Wildfire Automated Biomass Burning Algorithm

VIIRS	Visible Infrared Imager and Radiometer Suite
-------	--

1. Introduction and background

Wildland fire response and management represent issues of growing global importance. In the last 15 years, the number of large wildfires (or simply fires, hereafter) and annual area burned, particularly in the western U.S., has increased markedly, resulting in significant threat to public safety and the environment. Incident response and management have caused critical budget impacts due to overwhelming costs of suppression. For example, in an average year during 2005-2014, wildfires consumed 2.7M hectares at a cost of \$1.6B (suppression only) to federal agencies (NIFC, 2014). Only ~1% of ignitions in the U.S. become large escaped fires, i.e. fires that have exceeded initial attack capabilities and expanded beyond 40 ha of forest or 120 ha of shrub/grass (QFR, 2014). However, these fires have highest risk potential to firefighter safety and are responsible for most of the total burned area and suppression costs (e.g. NICC, 2013). The suppression costs form only a small fraction of the total societal losses from large wildfires that include loss of life and property and impacts on public health, economic activity, and environment (QFR, 2014). Consequently, rapid and prioritized response to fire ignitions that have a great risk to become large incidents could lead to high benefits to society.

Timely and informed management decision making critically depends on how quickly ignitions are identified and confirmed. Earlier detection often leads to a smaller fire size at initial attack, thus increasing the probability of containment (Hirsch et al. 1998). How rapidly the value of wildfire detection information decreases with time depends on various factors, including a human component. The authors are not aware of studies providing a quantitative account of this

issue. Fire managers and first responders are convinced that to contain potentially damaging wildfires, ignitions should be identified within the first hour, but preferably within minutes.

New ignitions over the continental U.S. are identified primarily by human observations, i.e. the general public, commercial airline flights, fire lookout stations, aerial reconnaissance during periods of high fire danger or ignition potential. Most of the ignitions are rapidly seen and reported. However, as the conventional discovery methods are non-systematic, infrequent, and/or geographically localized, there are routinely situations where a fire went undetected for hours or days, both in remote and populated areas (e.g. downed power lines in the overnight hours, smoldering ignitions after lightning events, incompletely extinguished or illegal campfires, etc.). Furthermore, after an initial report, significant confusion and uncertainty often remain about the incident location, magnitude, or its very existence, making it more difficult for first responders and managers to develop and execute an appropriate response strategy.

Under these circumstances, the thermal infrared (TIR) observations from currently operational environmental and weather satellite programs, such as NASA's Earth Observing System, NOAA's Polar-orbiting Operational Environmental Satellites (POES), its successor JPSS (Joint Polar Satellite System), and Geostationary Operational Environmental Satellites (GOES), have been considered as potential means to rapidly detect wildfire starts over large areas and be used for initially alarming or as a necessary confirmation of recent alarms received from conventional sources. Indeed, as these programs were launched to support a broad range of civilian applications, they offer a range of valuable practical advantages, including low per-application cost, global systematic coverage, operational stability, and long-term continuity.

Nevertheless, while active fires have been successfully mapped by these programs for decades (Justice et al. 2011; Csiszar et al. 2014; Prins et al. 2001, Schmidt & Prins 2003; Prins et al. 2010), the corresponding fire detection products have not significantly reduced the time to first detection of new ignitions (Martell, 2015). Measurements by polar-orbiting sensors are a few hours apart, often have significant data dissemination lags, and therefore they currently have a marginal value as early warning tools. Images from GOES do have sufficiently frequent temporal coverage of the Western hemisphere: normally, at 15-min time steps, and every 5-7 minutes under GOES Rapid Scan operations. However, they also have a coarse spatial resolution (e.g. $\sim 25 \text{ km}^2$ over California). Consequently, small-magnitude thermal anomalies at the pixel level during early phases of burning can be difficult to automatically discern from naturally dynamic background. Despite this and other factors complicating geostationary detection (Schmidt et al. 2012), previous individual case studies (e.g. Feltz et al. 2003, Weaver et al. 2004, Koltunov et al. 2012a) indicated that the high temporal coverage of GOES imagery could often be sufficient to provide early alarms about new ignitions. Thus, it is natural to ask a question: is the early warning potential of the GOES satellites already fully utilized by the current operational wildfire algorithm?

1.1 Early detection of new ignitions is a new type of satellite wildfire remote sensing

Wildfire remote sensing from GOES is operationally realized by the Wildfire Automated Biomass Burning Algorithm (WF-ABBA, Prins & Menzel 1992, 1994; Prins et al. 1998, 2001, 2003) that in the early 1990s pioneered geostationary wildfire remote sensing and recently

expanded to other geostationary satellites across the globe (Prins et al., 2010). The WF-ABBA algorithm was not specifically designed as an early warning tool; and its primary applications include fire weather analysis and forecasting; climate, land-use, and land-cover change research; emissions, aerosol, and trace gas modeling, and other environmental applications. Consistently with these applications, WF-ABBA was developed and optimized for the performance measures based on counting correctly classified pixels (i.e. pixelwise false positive and false negative rates) and maximizing the number of *eventually* detected incidents (Koltunov et al. 2012a). In contrast, early fire detection (EFD) systems need to inspect images for a very different type of targets: previously undetected ignition events that may span multiple pixels in GOES images (see Table 1). Furthermore, the primary objective of an EFD system is to detect new events as rapidly as possible, which is a low priority for most WF-ABBA users. Indeed, for a typical seven-day wildfire incident, a two-hour delay in initial detection in GOES data increases the pixelwise false-negative rate by about 1% (roughly, the ratio of delay time to burning time), with a similar expected effect on products like estimates of total gas emissions from this incident. However, such a delay is likely to greatly reduce the value of the detection information for initiating a timely tactical response.

Table 1: Primary objectives and features for two distinct types of geostationary wildfire remote sensing: Active Fire Monitoring and Early Fire Detection.

ACTIVE FIRE MONITORING e.g. WF-ABBA	EARLY FIRE DETECTION e.g. GOES-EFD
Maximize detected fire pixels	Maximize detected fire ignitions (incidents)
Minimize false detection fire pixels	Minimize false alarms

Estimate fire characteristics (radiative power, area, temperature)	Minimize time to initial detection
Globally, not regionally calibrated	Regionally and seasonally calibrated
Global coverage is essential	Deployment in selected regions, as needed
Operational system available	In research and development

109

110

111 Target objects of an EFD algorithm are by orders of magnitude more rare than fire pixels.

112 Although the incident in our example is present in as many as 672 GOES images (95 images per

113 day, with routine scanning), there is only one true target *object* for an EFD system for this

114 incident, whereas there can be nearly 2000 true wildfire *pixels* to detect (assuming without loss

115 of generality 3 fire pixels per image on average). Thus development and real-data validation

116 experiments may not have many true-positive examples to work with, unless the image

117 sequence is very large. This situation is further complicated, as higher-resolution imagery that is

118 effective at validating expected pixelwise performance (Schroeder et al. 2008a, 2008b) is too

119 infrequent to resolve ignition times, making the EFD developer rely on an often incomplete and

120 occasionally inaccurate wildfire report data to evaluate detection timeliness for incidents

121 (Koltunov et al. 2012a).

122

123 Furthermore, hot spot pixels from the same wildfire incidents do not occur at random locations

124 or random times: they have very strong temporal and spatial autocorrelations. Hence, to achieve

125 the same relative detection accuracy (i.e. # true positives / # positives) for incidents as for pixels

126 the false positive pixels would have to show a similar strength of spatio-temporal

autocorrelation as true wildfires, which would make it possible to report one false incident per hundreds or thousands of these pixels. Unfortunately, spatio-temporal autocorrelation between non-fire hot spot pixels is not strong, i.e. they appear more randomly in the image sequence. As a result, an algorithm that is reliable when target objects are *pixels* inevitably becomes far less reliable when the target objects are wildfire *incidents*. A practical illustration of this phenomenon was provided by Koltunov et al. 2012a, who evaluated WF-ABBA over California fire season 2006 and found that more than 75% of detected fire pixels were true wildfire pixels, but only between 17 and 40% of apparent new fire starts were true wildfire incidents. The challenge in keeping false positives under control is magnified, as the EFD goals discussed above necessitate detection of significantly subtler anomalies to prevent delays. Indeed, under Gaussian noise assumption the false positive probability grows exponentially with detection sensitivity; so, for example, detection with a Z-score (standardized residual) threshold of 1.75 instead of 3.5 entails a nearly 170-fold increase in the number of potential false positives.

Thus, monitoring/characterization of active fires and early detection of fire starts represent two clearly distinct types of wildfire remote sensing, although the same general term “fire detection” has been commonly applied to both. Table 1 briefly summarizes the differences between these two problems.

1.2 Toward GOES early wildfire detection capabilities

As of today, there is no GOES image processing system with a dedicated focus on maximizing detection timeliness while minimizing false initial alarms. What aspects of the GOES image analysis need to be improved to advance toward these goals?

The physical basis for TIR fire detection at subpixel scale (Matson & Dozier 1981) stems from Planck's law and underpins most of today's operational wildfire monitoring algorithms, including WF-ABBA, the MODIS Fire and Thermal Anomalies algorithm (Giglio et al. 2003; Giglio 2010), and their successors (Schmidt et al. 2012, Schroeder et al. 2014) adapted for the GOES-R Advanced Baseline Imager (ABI) and Visible Infrared Imaging Radiometer Suite (VIIRS) sensors, respectively. Conceptually, subpixel fire-candidates are found by detecting anomalously large brightness temperatures (BT) in a shortwave TIR (often called middle infrared) band centered near 4 μm that co-occur with large differences between the BT in the shortwave and a longwave TIR (e.g. $\sim 11\text{-}\mu\text{m}$) band.

While improvement in this physical principle is strictly constrained by the available satellite bands, the way the available bands are analyzed leaves much greater room for detection timeliness optimizations. All operational fire monitoring algorithms are implemented as several contextual and fixed-threshold TIR tests that are logically merged and coupled with auxiliary techniques and data along with the visible channel to filter out false positives. Thus, to detect anomalies these algorithms utilize multispectral information from only the *inspection* frame (in this paper, the term *frame* stands for one of the many multispectral images composing a satellite image sequence; and *inspection* refers to automatic analysis to detect and categorize anomalies in the image). The temporal dimension is only used (e.g. by WF-ABBA) as an optional post-processing step to verify consistency of single-frame detections, not for seeking fire-candidates as a class of temporal changes in the pixel values, which indicates an underutilized potential for improvement. This potential was illustrated by Koltunov and Ustin (2007), who applied a

multitemporal anomaly detection method called Dynamic Detection Model (Koltunov et al. 2009) to MODIS imagery and found that temporal dimension of satellite data could be significantly more informative than the spatial dimension for detecting thermal anomalies, such as fires. In line with those findings, in recent years, there has been a growing consensus about the need to incorporate a temporal change analysis and several studies suggesting improvements in active fire monitoring (Calle et al. 2006, Mazzeo et al. 2007; Schroeder et al. 2008a, Xu et al. 2010, Roberts & Wooster 2014). As a result, new algorithms have emerged (e.g. Mazzeo et al. 2007, Xu et al. 2010, Roberts & Wooster, 2014) offering different implementations of the multitemporal approach. However, the temporal domain analysis also introduces problems leading to new types of false alarms. Indeed, some dynamic events and changes in the image or in the monitored scene can be misinterpreted as fire ignitions. These include apparent changes in the pixel intensities due to image misregistration across time, cloud motion, intermittent sun glints, and other dynamic events and factors. As a result, development of a functionally complete and reliable early fire detection algorithm requires detailed analysis of the apparent anomalies and fairly complex inter-related additional processing steps. These complexities are often magnified, when the objective is to maximize the timeliness of initial wildfire detection, due to the need to elevate algorithm sensitivity and because retrospective analysis delays detection.

1.3 Paper objective and outline

To address the urgent need for timely and cost-effective information about new fire ignitions, and the limitations of existing operational satellite fire detection algorithms to provide early warning capabilities, the *Center for Spatial Technologies and Remote Sensing* (CSTARS) at University of California, Davis in collaboration with USDA Forest Service *Remote Sensing*

Applications Center (RSAC) have been developing a new GOES Early Fire Detection (GOES-EFD) system that specifically focuses on the timeliness and reliability of the first detection of new wildfire incidents. This paper presents the GOES Early Fire Detection algorithm version 0.4. In the following, we will use the acronym GOES-EFD 0.4 or simply GOES-EFD, where appropriate. With our algorithm, we evaluate the potential of alternative image analysis methodologies to identify new ignitions before operational capabilities realized by WF-ABBA 6.1 and before conventional wildfire identification mechanisms. In section 2, we begin with a conceptual overview of the system and then discuss individual components in sect. 2.3-2.7. An experimental test of the GOES-EFD using California 2006 fire season data is presented in section 3, followed by discussion and conclusion. A short preliminary version of this paper can be found in Koltunov et al. 2012b.

2. GOES-EFD system version 0.4

The development of the GOES-EFD algorithm benefits from years of research by different teams in the areas of subpixel fire detection physics (Matson & Dozier, 1981), multitemporal thermal anomaly detection (Koltunov & Ustin, 2007, Koltunov et al. 2009), contextual hot spot detection (Prins & Menzel, 1994; Giglio et al., 2003; Giglio, 2010; Schroeder et al. 2014), automatic image registration (Irani, 2002), and wildfire incident analysis (Koltunov et al. 2012a).

2.1 Major procedural components

The GOES-EFD system version 0.4 utilizes only two thermal bands of the GOES Imager: centered at $\sim 4\mu\text{m}$ and $\sim 11\mu\text{m}$, respectively. The use of the visible channel is planned for a

future version. Let BT_4 and BT_{11} denote the brightness temperature bands derived from these two original GOES bands; and the difference band $\Delta BT \stackrel{\text{def}}{=} BT_4 - BT_{11}$.

System operation includes several major procedural components that are executed in different combinations and variants. These components include:

1. Initial frame preprocessing (Fig. 1), including GVAR-to-BT conversion, single-frame cloud detection (sect. 2.4.1), and band-to-band registration (Koltunov et al. 2012b);
2. Image registration (Koltunov et al. 2012b);
3. Anomaly detection (sect 2.3), including multitemporal and contextual tests;
4. Cloud masking (sect. 2.4);
5. Pixel classification (sect. 2.5);
6. Temporal filtering (sect. 2.6);
7. Event tracking (sect.2.7).

Figure 1. Schema of the GOES-EFD preprocessing module

2.2 Operation overview

We define the term *scene* as a monitored region geographically corresponding to a rectangular subset of the GOES full scan image (an example is shown in Fig. 12). In an operational environment, GOES-EFD can process large monitoring areas (e.g. Western U.S. or the entire GOES Imager scan area) as a single scene. Alternatively, large areas can be divided into smaller scenes to be individually processed by separate instances of the GOES-EFD system software. The distributed processing by smaller scenes can lead to improved detection performance, owing to the reduced environmental and scan angle variability within the scene and because the processing can begin sooner. Among potential downsides are the overhead costs of operation and possible shortage of useable pixels when the scenes are too small. Optimizing the partition and operational coverage of large areas by GOES-EFD scenes is an interesting research topic outside the scope of this paper.

For a given scene, GOES-EFD operates in two modes: Training mode and Detection mode. The Detection mode includes all the above seven components and refers to a regular mode of operation, in which GOES multispectral frames are automatically inspected to detect fires. The Training mode includes only the first four procedural components.

Figure 2. Schema of the GOES-EFD Training stage operation (see text in sect. 2.2 for details).

The Training stage (Fig. 2) is normally performed for each monitored scene once a year, before the beginning of the Detection mode, e.g. on the onset of the fire season. In the Training mode, GOES-EFD prepares several static datasets and calculates static parameters, including:

- a coordinate system for the scene, which can be a coordinate system of e.g. any specific GOES frame and which acts as an intermediary parameter facilitating image alignment and georeferencing;

- a set of GOES images defining a static (i.e. fixed) baseline period and representing background (i.e. “no anomaly”) conditions on the ground for anomaly detection during the Detection stage, as further detailed in sect. 2.3; these images are called *basis images* (sect. 2.3.1).
- registration (i.e. relative image motion) parameters for the basis images; and
- cloud/missing-value masks for the basis images.

Toward these goals, the algorithm analyzes a large sample (e.g. hundreds) of past GOES images (Fig. 2) of the scene, which we term the *training* frames. The training frames should represent diurnal and preferably also seasonal thermal variation, e.g. all images from previous 1-2 months or last year, as available. High cloud cover scenes may require more training frames. First, all training frames are preprocessed (Fig. 1) and automatically registered toward a number of reference frame-candidates (Koltunov et al. 2012b). If the average estimated pixel motion is less than 3 pixels (or another *a priori* known bound) and the R^2 of the linear model between the reference images and the image being registered exceeds 0.3, then registration is considered successful. Next, the reference image set that provided successful registration of the largest number of frames is automatically selected. Because the outcome of this selection process naturally includes the estimated image motion parameters for all training images, a multitemporal cloud detection algorithm can be applied, resulting in refined cloud masks for all training images. This algorithm is called the Retrospective Cloud Detection (RCD) algorithm and described in sect. 2.6.3. Finally, well-aligned frames with low cloud cover (e.g. <5%) representing the 24-hour cycle at a half-hourly to hourly step are sorted by the increasing cloud cover; and for each half-hour/hour band the least cloudy images from different dates are automatically selected as static basis image candidates. As a recommended optional step, these

images can be then interactively reviewed, in particular, to make sure that the undetected clouds, residual misalignment, and other possible artifacts are minimal. The presence of potential active fires is not a factor in approving or rejecting a basis image. A rejected basis image candidate can be replaced with an alternative training image from a similar observation hour on that or a different day, as available. In general, more basis images tend to provide a better detection performance at the expense of computation time.

Figure 3. Schema of the GOES-EFD Detection stage operation (see text in sect. 2.2 for details)

During the Detection stage (Fig. 3), the current image is first preprocessed (Fig. 1), and the image motion parameters are computed. Next, the Iterative Anomaly Detection and Classification (IADC) module (sect. 2.5) is applied to assign pixels into one of 12 thematic classes. Each iteration of IADC performs anomaly detection tests and combines their outputs with auxiliary information, such as land and ocean masks. The resulting thematic classes include: "no-anomaly", clouds, several fire confidence classes related to the magnitude of the thermal anomaly, and other classes. Frames with more than 90% cloud cover over land or those for which initial preprocessing failed for any reason are excluded from further analysis. The classification output product from IADC can be immediately delivered to the user. However, to reduce false positives, GOES-EFD employs a temporal filter (sect 2.6), which rejects fire pixels that were not detected in the previous frame. The filtered fire detections at the chosen level of confidence are further analyzed across spatial and temporal dimensions by the *Event Tracker* algorithm (sect. 2.7) to extract *events*, i.e. groups of confidently detected pixels that are likely to represent the same incident. The new, previously undetected events extracted by the Event Tracker are the primary output of the GOES-EFD algorithm. In the following sections we discuss the individual components and modules of the system and their use during the Training mode and Detection mode in more detail.

2.3 Anomaly Detection

2.3.1 Multitemporal anomaly detection with Dynamic Detection Model

The GOES-EFD module for detecting anomalous temporal changes in the image sequence is a core capability facilitating image registration, cloud detection, and fire pixel detection. It is based on the Dynamic Detection Model (DDM, Koltunov et al. 2009), which can be viewed as a

special type of multiple regression methods. This model stems from a generic physical hypothesis that temporal and spatial invariants (separable parameters) play a dominant role in forming many remote sensing image datasets, including for instance, images of radiance, brightness temperature, vegetation indexes, to name only a few. Therefore, at a detection time t and pixel location s , a DDM models the background (“no anomaly”) value $W(s, t)$ as a space-invariant dynamic combination of values that were observed at that pixel previously, at some number P of previous time moments termed *basis* times and denoted $t_1, \dots, t_P$. As discussed by Koltunov et al. 2009, the basis times define both the baseline time period for comparison and determine the range of surface, sensor, and environmental parameters in the scene that DDM models. As we further detail in sect. 2.3.3, GOES-EFD uses two linear models of background:

$$W(s, t) = \beta(t)_0 + \beta_1(t)W(s, t_1) + \dots + \beta_P(t)W(s, t_P). \quad (1)$$

Both models of type (1) are applied to each of the three brightness temperature bands defined in sect. 2.1, and thus the generic variable W we use may stand for either BT_4 , or BT_{11} , or ΔBT . The unknown parameter vector $\beta(t)$ depends on the observation conditions at time t but is the same for all pixel locations s . Thus in mathematical terms, model (1) projects the inspection image band, as a single high-dimensional vector, onto a linear non-orthogonal vector subspace spanned by the images $W(s, t_1), \dots, W(s, t_P)$ and a vector of all ones. These images are called *basis* images because they constitute a basis in the above subspace. To reconstruct the background, the unknown coefficient vector $\beta(t)$ is estimated in the least-squares sense using a random set of training pixels that are not *a priori* known anomalies. Pixels that are outliers with respect to the prediction model (1) are considered anomalies. Using a physical approach, Koltunov et al. (2009) derived the ranges of weather and sensor characteristics and the surface thermophysical properties, under which radiance can be accurately modeled by the linear model

(1) with eight basis images. In general, more complex practical scenarios can be modeled by increasing the number of basis times, using a non-linear background model, or both (Koltunov & Ustin 2007), at the cost of extra computational power.

We see from the discussion in the previous paragraph that the term “multitemporal” that we use to characterize the DDM-based background modeling should not be interpreted as if only the temporal dimension from the pixel in question is used to estimate $W(s, t)$. Perhaps, a more precise (though less concise) term would be “blended spatial-multitemporal”. This is because indirectly through $\{\beta_k(t)\}_{k=0}^P$, the predicted value depends on the dynamically chosen training pixel values at basis times and at the inspection time. This differentiates DDM from the time series analysis techniques in which each pixel is processed independently across temporal domain (e.g. Verbesselt et al. 2010, Roberts & Wooster 2014).

A characteristic property of DDM is that the time variable t is not directly used as a predictor variable. Therefore, the image on the left-hand side of (1) can be a past image modeled using chronologically subsequent basis images on the right-hand side of (1). This property is utilized by GOES-EFD for retrospective cloud detection (sect. 2.4.3) during the Training stage.

The forward application of the model (1) to a given pixel requires that all basis image values at that pixel correspond to the same object in the scene. Therefore, a single set of basis images can meet this requirement only if all basis images are completely free of clouds, fires, and other anomalies. This is not a realistic scenario, unless the number of basis images in the scene-wide model is small. Therefore, to avoid unnecessary underfitting of the background model (1) by

using too few basis images, we analyze missing value (i.e. “known anomaly”) masks for various subsets of the basis images $t_1, \dots, t_p$ and determine which pixels can be modeled for each subset. For example, pixels with missing values only at t_1 cannot be modeled using the entire set $t_1, \dots, t_p$; yet they can be modeled using $t_2, \dots, t_p$ or any subset thereof, perhaps with larger error. Pixels with missing values only at t_1 and t_2 can be modeled using $t_3, \dots, t_p$ or any subset thereof, and so on. Using this information, we compile an ordered list of basis time subsets by removing basis times one-by-one, until the current list of subsets can model all but an acceptable number of land pixels (according to the auxiliary land mask), or until there is only one basis time in a subset. In this way, we can avoid the exhaustive search and the associated combinatorial computational complexity.

Next, for each subset in the list of basis time subsets, a background model is estimated by robust stepwise regression using training pixels chosen at random from where the model is applicable, i.e. where none of basis images defining the model have a known anomaly, as discussed in the previous paragraph. As a result, multiple DDM-based background models are constructed for each pixel, thus providing a pool of options for background estimation. GOES-EFD 0.4 uses only one option per pixel, which is the model with the maximal adjusted determination coefficient. A typical example of scene coverage with 32 multitemporal models per pixel is given in Figure 4c), where the scene is shown in Figure 12 and referred to as Scene 1, hereafter (further details on the fire season 2006 experiment datasets are available in section 3.1). In Figure 4c), different models are represented with different colors, and the corresponding model r.m.s.e. map is displayed in Figure 4d). The dark blue areas in these figures represent clouds detected by the GOES-EFD cloud detection algorithms (sect. 2.4 and Fig. 7) and also other

known anomalies and non-burnable surfaces, such as water. As seen from Figure 4c), for most pixels, the background can be modeled with a single set of basis images, in this case model #13 with 28 out of 29 basis images and r.m.s.e. of $\sim 0.4^\circ\text{K}$. Areas, for which models with fewer basis images are used due to missing values at basis times, such as model #2 and #32, typically have greater r.m.s.e. of background estimation.

Figure 4. Example of thermal background modeling with 32 multitemporal DDM-type models (see also text in sect. 2.3.1) for a scene in central California shown in Fig. 12. GOES input brightness temperature bands, BT_4 and BT_{11} (sect. 2.1), are shown in the top row images. For each pixel, GOES-EFD analyzes a pool of applicable DDM-models and estimates the best DDM-model shown in image (c). The absolute accuracy (r.m.s.e.) of background modeling with different models is shown in image (d). The cloud areas masked with dark blue color in (c) and (d) were detected by the Operational Cloud Masking (OSM) algorithm (sect. 2.4.2 and Fig. 7). The frame timestamp is 2006.08.03-19:00 PST.

Detection of subtle subpixel anomalies using multitemporal background reconstruction requires subpixel level co-registration of frames. The GOES image navigation and registration (INR) system is designed to dynamically adjust GOES Imager scan mirror pointing (Carr 2009). However, even under normal operation of onboard systems, minor inaccuracies in INR may lead to temporal misalignments in TIR bands, which may reach half a pixel or more (Boeing, 2006), in addition to regularly occurring band-to-band misregistration of a smaller magnitude (Grotenhuis et al. 2012). In addition, there have been several instances of a major failure of the INR components when image misalignment reached multi-pixel scales, including GOES-10 in 2006 and most recently, GOES-15 star tracker failure in April 2015. To mitigate these issues, GOES-EFD includes an automatic thermal image registration method (Koltunov et al. 2012b) that is a member of a family of direct methods (Hanna, 1991, Irani, 2002) for image alignment or motion estimation. The anomaly detection is performed in the inspection image coordinate system, i.e. the basis images are warped (moved and interpolated) toward the inspection image before substitution into (1). The binary masks are also aligned with the inspection image. The GOES-R ABI images are planned and expected to have significantly better alignment than the current generation of GOES; and they will be registered and resampled on the ground before distribution to the users as a Level 1B product. The magnitude and spatial and temporal distribution of the navigation and registration errors in these images is an important question to be answered after the GOES-R launch.

2.3.2 Contextual anomaly detection

Decades of operational success in contextual active fire detection have clearly demonstrated that that the spatial context should not be dismissed as an ignition information source. The contextual anomaly (“hot spot”) detection in GOES-EFD is based on the same principle as the

WF-ABBA and the MODIS Fire and Thermal Anomalies algorithm. The background value at a pixel is estimated by averaging values for valid neighboring pixels in a window. The window size is increased starting from 3x3 pixels until eight valid neighbors are available or the valid neighbor fraction in the window exceeds 0.25. Windows larger than 11x11 pixels are not considered. There are important differences between our implementation of this idea and the other wildfire algorithms. Our contextual analysis is applied to brightness temperatures, does not utilize external surface emissivity data, and does not directly account for oversampling along the GOES Imager scan. Also, the rule for determining valid neighbors is different: in GOES-EFD, valid pixels are determined based on classification of the multitemporal anomaly detection outputs (discussed in sect. 2.5), i.e. entirely independent of the contextual test itself. This indirectly mitigates the influence of anomalies on the background estimation due to oversampling.

2.3.3 Anomaly detectors

Using multitemporal and contextual approaches to anomaly detection, three anomaly detection tests (for brevity called *detectors*, hereafter) are applied for each BT band:

- DDM1: a multitemporal detector using the same static basis images (chosen in advance, during the Training stage, as discussed in sect. 2.2) for predicting background of all inspection images.
- DDM2: a multitemporal detector using the basis images of DDM1 and an image taken at $t-30$ min. If the $t-30$ min image is not available or if it is the frame immediately before the inspection frame, then an earlier image is used.
- CNTXT: a contextual hot spot detector.

Each detector is applied separately to each of the brightness temperature bands, BT_4 , BT_{11} , and ΔBT . The background prediction residuals are standardized by dividing by their respective r.m.s.e. estimates, σ :

$$Z = \frac{BT_{\text{observed}} - BT_{\text{predicted}}}{\sigma}.$$

These standardized residuals, also called *Z-scores*, are used to detect and classify changes in the scene, as described in the following sections. In the following, we will be using subscripted variables: Z_4 , Z_{11} , and Z_{Δ} , to indicate the corresponding BT band. Also when necessary, we will be labeling the Z-scores with a superscript, such as “ddm1”, “ddm2”, or “cntxt”, specifying the corresponding anomaly detector. Omission of a superscript should be interpreted as “for any detector”, unless explicitly stated otherwise or obvious from the context.

As seen from the detector definitions, they analyze images from different and largely complementary perspectives. The contextual approach has been extensively discussed in the past (see e.g. Giglio 2003, Schmidt et al. 2012), and so we focus our discussion primarily on multitemporal detectors. The multitemporal detectors differ by their respective sets of basis images, and thus they are sensitive to different temporal scales, types, and history of change in the scene.

Because DDM uses a fixed baseline, the background reconstruction quality for frame $W(t)$ depends on the generalization capability of the basis images and the change status at time t , yet it is completely independent of the thermal signal history or history of detection before time t . In this way, the static-baseline DDM1 is a system with no short-term memory. In other words, DDM1 provides frame-wise snapshots for *cumulative effects* of scene changes occurred since

the end of the baseline period. Therefore, it is effective at consistently re-detecting ongoing wildfire events, persistent clouds, and other persistent changes in land cover. It also performs consistently following extended periods of cloudiness and after protracted data gaps in GOES image sequences. Furthermore, no matter how slowly an anomaly intensifies, it will be detected by DDM1, as soon as the cumulative effect of this event in the BT reaches a detectability threshold. On the other hand, even high magnitude change events can go undetected, if the total cumulative effect is small, as is in the case of repeat land cover changes. Another weakness of DDM1 is that after a permanent land cover change, such as burn or urban development, it can produce persistent anomalies unrelated to new wildfire ignitions. These types of changes typically affect a very small fraction of the scene pixels, thus minimizing the problem.

In contrast, DDM2 has the capability to respond to changes relative to *recent* baseline conditions. Such changes are often associated with new fire ignitions. Slowly progressing changes, however, may go undetected by DDM2, especially when the image frequency is high, e.g. during the GOES Rapid or Super Rapid Scan mode. For pixels that were not anomalies in the basis images, background estimation is usually significantly more accurate with DDM2 than with DDM1, because just one basis image $W(t-30 \text{ min})$ at these pixels already provides a close approximation to the inspection image $W(t)$. In this case, the information in the other (static) basis images, in effect, is used mainly to approximate the thermal gradient $\Delta W = W(t) - W(t-30)$; whereas in case of DDM1, the same static images model both, $W(t-30)$ and ΔW . The detector accuracies are illustrated in Figure 5a) showing BT_4 reconstruction r.m.s.e for approximately an 11-day period in August 2006 at a GOES pixel, which seven years later would include the origin of the 3rd largest wildfire in California history – the Rim Fire. Figure 5b-c)

show typical examples of, respectively, daytime and nighttime BT_4 approximation with these detectors for that pixel. It can also be seen from Figure 5 that in absolute units, all three detectors are more accurate at night time, when spatial and temporal gradients of brightness temperatures are minimal. At night the contextual detector tends to outperform DDM1. Overall, the most sensitive detector is DDM2, followed by DDM1, and CNTXT.

Figure 5. Performance of GOES-EFD detectors. (a): The r.m.s.e. time series from 2006.08.04-0300 to 2006.08.15-0930 PST for the three detectors used by GOES-EFD: DDM1, DDM2, and CNTXT (defined in sect.2.3.3), applied to model the observed GOES-11 3.9 μm brightness temperatures (BT_4) during California 2006 experiment (section 3). (b) and (c): typical examples of day time and night time reconstruction of GOES-11 BT_4 values in August 2006 by the above three detectors for a non-fire pixel location with no cloud cover observed or detected. The pixel GOES-11 coordinates (line, element) = (4148, 17472) on 2006.08.03-1900 PST; (37.83N, 120.10W).

An example in Figure 6 illustrates how “labor sharing” between the two DDM-based detectors works when a change event, such as a fire, starts at time t_n . This figure displays a 15-min step time series for the observed BT_4 and the BT_4 approximated by the two multitemporal detectors at a GOES pixel s , where an anomaly from an actual fire first appears around 16:45 local time. In general, the more sensitive DDM2 (dark red crosses in Fig. 6) statistically tends to provide closer approximation for background and stronger response to change events than DDM1 (light blue squares), and thus the fire has a greater chance to get detected for the first time. Yet regardless of whether the anomaly in pixel s gets detected or missed in frame $W(t_n)$, for the next frame $W(t_{n+1})$ the DDM2 does not use $W(t_n)$ as a basis image. In this way, we avoid violation of the “no-anomaly” assumption for basis image values and reduction in the chance for the second detection at $W(t_{n+1})$ by DDM2. In the following frame at time t_{n+2} , image $W(t_n)$ normally becomes a basis image for DDM2 for pixel s , unless an anomaly was detected at (s, t_n) and flagged as a missing value (brown stars in Fig. 6), in which case DDM2 reduces to DDM1 (sect 2.3.1). For example, in Figure 6, two frames after the initial anomaly detection at 16:45, the BT_4 predictions for both models become identical until 18:15, when the 17:45 basis image for DDM2 no longer has a known anomaly. Thus, the static DDM1 takes over when DDM2 is deemed unreliable and provides detection as long as the cumulative effect signal is significant. If the anomaly detected at time t_n was a fire, then by t_{n+2} the fire has been burning for at least 45 min under GOES Regular Scan (25 min under Rapid Scan), and therefore it often presents a sufficiently strong signal to be re-detected by the less sensitive DDM1 model.

Figure 6. GOES-11 15-min step time series illustrating brightness temperature approximation by detectors DDM1 and DDM2 (as discussed in sect. 2.3.3), following a detected fire at ~16:45 PST at this pixel. The detected fire points in the BT time series are marked with dark brown stars. For the first two detections (at 16:45 and 17:00), the DDM1 Z-scores (2.9 and 4.5) are lower than those of DDM2 (4.4 and 8.3). In images from 17:15, 17:30, 17:45, and 18:00, DDM2 reduces to DDM1, due to an anomaly in the corresponding image from 30 min ago. The pixel GOES-11 coordinates (line, element) = (4332, 17772) on 2006.08.13-1700 PST; (35.58N, 118.52W).

Also, these detectors have different practical application constraints. In particular, because baseline images of DDM1 are fixed, GOES-EFD prepares them in advance during the Training stage, thus making it feasible to apply retrospective automated analyses or interactively correct or replace some of the basis images. As a result, the influence of clouds, fires, and other atypical objects in the static basis images that violate the background model assumptions can be minimized. In contrast, the dynamic detector DDM2 can rely only on the less accurate fully automatic real-time techniques to determine pixels where it cannot be applied. These techniques include:

- 1) GOES-EFD real-time cloud, fire, and other anomaly detection outputs for the $t-30$ min basis frame;

2) a robust stepwise regression that can either exclude the $t-30$ min frame from the model, if that increases the adjusted determination coefficient, or flag these pixels as outliers to the regression model, or both.

2.4 Cloud Detection and Masking

Clouds markedly affect fire detection, rendering cloud delineation indispensable for a fire detection system. Cloud cover information is required during different stages of the system operation, with their respective constraints and quality requirements. Therefore, depending on specific circumstances and a processing stage, GOES-EFD can use the following three cloud detection algorithms:

- 1) Single-Frame Cloud Detection (SCD),
- 2) Operational Cloud Masking (OCM), or
- 3) Retrospective Cloud Detection (RCD).

The algorithms in this list are ordered by the increasing complexity and are discussed below.

2.4.1 Single-Frame Cloud Detection (SCD)

As the name suggests, this algorithm uses only information data from the inspection frame and conservatively (i.e. with a low false positive rate) detects opaque high altitude clouds or snow. Presently, it uses a simple condition $BT_{11} < 270K$ that may need to be adjusted depending on the geographic region and monitoring seasons. The SCD algorithm is used as a pre-screening procedure during both Training and Detection stage. In particular, it is applied before a more sophisticated multitemporal technique, and when a multitemporal cloud detection algorithm cannot be applied. Also, single-frame clouds define regions that cannot be used to estimate the frame or band-to-band motion parameters.

2.4.2 Operational Cloud Masking (OCM)

The OCM algorithm is a part of the Iterative Anomaly Detection and Classification block (sect. 2.5). It is applied during the Detection mode (Fig. 3), following frame-to-sequence registration. This algorithm combines Z-scores obtained from the static multitemporal detector DDM1, using an initial cloud mask and a land mask as ancillary inputs. The initial cloud mask can come from SCD, a previous iteration of OCM, or another source; and it is supposed to flag clouds conservatively, because these detections are immediately accepted by the OCM. For example, in Figure 7 the SCD detects the coldest clouds (shown in dark blue colors in Fig. 7c) that are excluded from modeling background with DDM1 and from image motion estimation. In the next step (Fig. 7d), the multitemporal OCM inherits these single-frame detections and masks out additional cloud and fog pixels colored in lighter shades of blue and corresponding to detections by tests C1-C4 described below.

Figure 7. Two-stage cloud masking illustration (see also text in sect. 2.4). GOES input brightness temperature bands, BT₄ and BT₁₁ (sect. 2.1), are shown in the top row images. The first stage applies SCD (Single-frame Cloud Detection) algorithm (c), whose detections are flagged as missing values for the multitemporal analysis using the Operational Cloud Masking algorithm and added to the final cloud map (d). The frame timestamp is 2006.08.03-19:00 PST.

The OCM algorithm flags a pixel as “cloud” if any of the tests C1-C5 below returns true for that pixel.

Test C1 detects cold high-altitude clouds or snow as an anomalous drop in BT₁₁:

$$Z_{11}^{ddm1} < -3.0 \text{ OR } (Z_{11}^{ddm1} < -2.0 \text{ AND } BT_{11} < 275). \quad (2)$$

Test C2 is looking for nighttime clouds and fog using the condition:

$$Z_{11}^{ddm1} < -1.5 \text{ AND } Z_{\Delta}^{ddm1} < -1.5. \quad (3)$$

Test C3 flags water body pixels in which DDM1 detected anomalous changes in BT₄, BT₁₁, or ΔBT :

$$|Z_4^{ddm1}| > 2.0 \text{ OR } |Z_{11}^{ddm1}| > 2.0 \text{ OR } |Z_{\Delta}^{ddm1}| > 2.0 \text{ OR } (|Z_4^{ddm1}| > 1.0 \text{ AND } |Z_{\Delta}^{ddm1}| > 1.0). \quad (4)$$

These changes can be due cloud, fog, surface reflections, other changes, and unchanged surfaces poorly modeled by DDM1 for any reason. The aggressive detection provided by this test intends to account for greater complexity of modeling background conditions over water and allow anomalies with lower confidence to be excluded from background estimation.

Test C4 attempts to identify thin or warm clouds or haze using the following condition:

$$(Z_{11}^{ddm1} < -2.0 \text{ AND } Z_{\Delta}^{ddm1} > 2.0) \text{ OR } \dots$$

$$(Z_4^{ddm1} < -1.0 \text{ AND } Z_{11}^{ddm1} < -1.5 \text{ AND } Z_{\Delta}^{ddm1} > 1.5) \text{ OR } \dots$$

$$(Z_4^{ddm1} < -2.5 \text{ AND } Z_{\Delta}^{ddm1} > 2.5). \quad (5)$$

2.4.3 Retrospective Cloud Detection (RCD)

The Retrospective Cloud Detection algorithm is applied only during the Training stage (sect. 2.2). This algorithm cannot be applied in the real-time mode, because for a given frame, it takes advantage of the entire input sequence, including subsequent frames. The RCD algorithm proceeds iteratively as follows:

- 1) Initially assign set S to be the entire sequence.
 - a. Define the set of basis images B as the least cloudy images from S at hourly steps.
- 2) Define the set of inspection images I as non-basis images in S .
- 3) Detect clouds in the inspection images:
 - a. Estimate background values in the current set of inspection images, using DDM1 with known clouds/anomaly masked out for the background estimation purposes.
 - b. Apply tests C1-C4 of the OCM algorithm.
- 4) Use the resulting cloud maps to detect clouds in the basis images:
 - a. Assign: $S \leftarrow I$, then $I \leftarrow B$, and repeat step 3).
- 5) Repeat steps 1)-4) using the updated cloud maps.

2.5 Iterative Anomaly Detection and Classification (IADC)

module

The IADC module is activated at the Detection stage after the inspection frame has been preprocessed (Fig. 3) and registered with respect to the scene coordinate system. We define a

626 fire-candidate (FC) test, which identifies a fire pixel-candidate at the confidence level γ as a
 627 pixel for which the condition

$$628 \quad Z_4 > \gamma \text{ AND } Z_{\Delta} > \gamma, \text{ for } \gamma > 2.0, \text{ where } \gamma = \min(Z_4, Z_{\Delta}) \quad (6)$$

629 holds for any multitemporal detector (DDM1 or DDM2). Based on their values of γ , such
 630 pixels are initially binned into fire confidence classes that for the sake of this presentation can
 631 be named: “very low” ($2 < \gamma \leq 2.5$), “low” ($2.5 < \gamma \leq 3$), “medium-low” ($3 < \gamma \leq 3.5$), “medium”
 632 ($3.5 < \gamma \leq 4$), and “high” confidence ($\gamma > 4$). GOES-EFD v.04 reports only medium and high
 633 confidence detections; lower confidence detections are utilized for filtering false positives.

Figure 8. Schema of the GOES-EFD Iterative Anomaly Detection and Classification (IADC) block (see text in section 2.5 for details).

The current version of IADC consists of two passes shown in Figure 8. The primary goal of Pass 1 is to identify pixels whose thermo-physical properties have significantly changed, e.g. clouds, higher intensity fires, and miscellaneous anomalies. These pixels will not be used during the Pass 2 for determining the coefficients $\beta(t)$ in the background model (1) or for

contextual background estimation. Pass 1 proceeds in three iterations or until the class membership does not change for more than 99.95% of land pixels, whichever comes first. In each of these iterations, the following steps are performed:

- apply detector DDM1 to BT_4 , BT_{11} , and ΔBT bands;
- identify potential fire pixels by the FC test (6) applied to DDM1 only;
- detect cloud pixels with the OCM algorithm (sect. 2.4.2).

Pass 2 begins with the following steps:

- apply detectors DDM1, DDM2, and CNTXT to bands BT_4 , BT_{11} , and ΔBT ; these detectors estimate background values using only land pixels classified as background by Pass 1;
- identify potential fire pixels by the FC test applied to DDM1 and DDM2 and use $\gamma = \max\{\gamma^{ddm1}, \gamma^{ddm2}\}$;
- detect cloud pixels with the OCM algorithm, which may reject some of the fire pixel candidates resulted from the FC test.

Pixels classified as fires based on temporal information as described above are subject to additional false-alarm elimination tests. These tests use the contextual detector Z-scores to detect possible hot spots by the following condition:

$$Z_{11}^{cntxt} > 1.5 \text{ AND } Z_{\Delta}^{cntxt} > 1.5. \quad (7)$$

A fire pixel candidate is eliminated if condition (7) is false (i.e., not a hot spot) and any of the following conditions returns true:

- the fire is detected by DDM2 and was an anomaly with $Z_4^{ddm2} < -2.0$ in the image at $t-30$ min;

- the fire pixel is detected by DDM2, not detected by DDM1, and was not assigned to any of the fire or cloud classes in the image at $t-30$ min;
- the fire pixel is adjacent to a detected cloud pixel in any of the frames time-stamped t , $t-30$ min, or $t-15$ min. Areas surrounding detected clouds often represent thin clouds, or locations of clouds in recent frames, and therefore such detections are more likely to be false positives;
- the fire pixel is not adjacent to a hot spot pixel defined by (7).

The first three tests are aimed at eliminating false positives sometimes committed by the multitemporal detector DDM2 mainly due to undetected clouds in the basis images or possible residual misregistration.

Furthermore, a fire pixel that passed through the above filters is rejected if:

- $BT_4 < 290K$, or
- $Z_{11}^{ddm1} < -2.0$.

The remaining fire pixels are subject to an object-based analysis aimed to find detections that form valid spatially connected regions. The Event Tracker module (described in detail later in sect. 2.7) will use this information to dismiss new detected fire-like events in the invalid regions. The validity of a region is defined in terms of the number of pixel-members, their anomaly confidence values, and the region orientation in the image coordinate system (see example in Fig. 9). A single-pixel fire region is considered valid, and medium or high confidence ($\gamma > 3.5$, see eq. (6)) fire regions with 5 or more pixels are considered invalid.

Remaining regions are flagged invalid if any of the following conditions holds true:

- The region has four or more fire pixels detected at any confidence ($\gamma > 2.0$) and no fire pixels detected with high confidence ($\gamma > 4.0$). Such regions often appear as uniformly

low magnitude anomalies spread over large areas. For example, in Figure 9b) all five pixels of the boxed object are detected at $2 < \gamma < 4$ (in this example Z_4 is equal to γ and shown as numbers inside the pixels). Therefore, although pixel with $Z_4=3.72$ is a medium-confidence fire-like anomaly, it is considered to be most likely caused by the same reason as the as the entire object and thus inconsistent with wildfire dynamics and the GOES pixel footprint size and the oversampling pattern.

- The medium or high confidence region ($\gamma > 3.5$) is vertically oriented in the image coordinate space. We fit an ellipse to the points representing centers of region pixels by computing the region central moments. If the angle between the ellipse's major axis and the x-axis of the image plane is between $-\pi/4$ and $\pi/4$, then the region is a valid region. The rejection of vertically oriented regions is based on the GOES Imager oversampling in the x-direction.

Although the object-based analysis can in some cases reject large ongoing true wildfire incidents, incidents adjacent to undetected reflective clouds, new ignitions a few kilometers apart from another wildfire, and some other rare types of new ignitions, it also helps reduce spatially contiguous false positives due to undetected clouds or fog, temporal misregistration, and error in background modeling by the DDM due to underfitting.

2.6 Temporal filter

The IADC output is the inspection image classification derived using information available at the time of inspection image acquisition. Although it is the most rapidly available fire product, it can also have large false positive rates. Thus, to reduce the number of false positives, a temporal filter is applied, masking out fire pixels until they get detected in two consecutive frames at a medium or high confidence ($\gamma > 3.5$). By delaying detection, the temporal filter can

give an order of magnitude reduction in false positive rates (Koltunov et al.2012b). To match fire pixels across time, we first account for pixel displacements due to frame-to-frame misalignment, and then apply a one-pixel buffer. This prevents the filter from missing repeat detections in cases when combination of nearest-neighbor interpolation and imperfect motion estimation results in a 1-pixel misalignment of the interpolated frame-wise anomaly masks. The filter of GOES-EFD 0.4 is suboptimal, particularly because it does not take into account the time interval between frames and gaps in clear pixel data due to cloud cover.

Figure 9. Fragments of DDM2 Z-score images with black dashed boxes indicating examples of (a) valid multi-pixel objects that are accepted as possible new fires by the IADC block (sect. 2.5) and (b) an invalid object where a new fire is rejected as a likely false positive. Anomalies known from IADC Pass 1 are marked with a NaN (not-a-number) symbol. See text in section 2.5 for details.

2.7 Event Tracker

Following Koltunov et al. 2012a, we introduce the term *detected efd-event* (also *efd-event*, or simply *event*, hereafter) as a set of GOES-EFD designated fire pixels in the image sequence that, according to the algorithm, should be considered by the user of the fire detection information as a possible single fire incident. An efd-event may or may not correspond to an actual wildfire. Similar to actual fire incidents, these detected events are dynamic objects in the GOES image sequence: they may persist in more than one frame and consist of multiple pixels per frame. By definition, every pixel that is flagged fire by the IADC is a member of exactly one efd-event.

As the GOES-EFD primary targets are new incidents, it includes an event analysis module, called *Event Tracker* (see the concept illustration in Fig. 10), which decides whether a fire pixel represents:

- *a new efd-event*, i.e. a group of pixels to be considered a previously undetected ignition,
- or
- *an ongoing efd-event*, i.e. an efd-event that was already detected by the algorithm in previous images and therefore assumed to have been reported to first responders.

Figure 10. A concept illustration for the GOES-EFD Event Tracker. The Event Tracker morphologically analyzes recent detection history to decide whether a fire pixel is a new event (a previously undetected possible ignition) or an ongoing event. New events are the primary output of GOES-EFD. See text in section 2.7 for details.

The method used by the Event Tracker is based on analyzing temporal evolution of spatially connected components (c.c.) formed by confidently detected fire pixels ($\gamma > 3.5$) in the GOES image coordinate system. We note that this temporal analysis complements the single-frame analysis of connected components by IADC (sect. 2.5). The algorithm is controlled by two parameters: h and b . The history length h is the size of a temporal window during which spatially overlapping (up to the buffer distance b) connected components are considered the same event. GOES-EFD 0.4 uses $h=48$ hours and buffer $b = 6.8$ km, which is approximately 1.2 times the linear ground size (in the north-south direction) of the largest GOES-West thermal infrared pixel in California. For a frame collected at detection time t , the Event Tracker combines fire pixels into efd-events as follows:

1. Group fire pixels into spatially connected components.

2. Initially, and also when there are no frames or detections available during last h hours, each c.c. is considered a new efd-event.

3. When past fire pixels are available, for each fire pixel (x, t) find the nearest location x_f that was flagged as a fire pixel at least once during the last h hours. If $\|x - x_f\| < b$, then pixel (x, t) is termed a *re-detected* fire pixel. Let $E(x_f)$ denote the event corresponding to x_f . Because $E(x_f)$ was detected before t , it is obviously an ongoing event at time t .

After the efd-events have been identified in the inspection frame, the algorithm performs steps 4 and 5 below to designate each event as a new event, or as an ongoing event. In Figure 11 exemplifying these steps, re-detected pixels are light grey and other pixels are dark grey. The label inside each pixel denotes the event to which the pixel will be assigned by steps 4 and 5.

4. Connected components with no re-detected pixels become new events (Fig. 11a).

5. For connected components that do include re-detected pixels (Fig. 11b-11h), the event assignment requires additional analysis:

a) First, each re-detected pixel-member (x, t) is assigned to its corresponding ongoing event $E(x_f)$ matched by step 3 (i.e. all light-grey pixels in Fig. 11b-11h are labeled with an ongoing efd-event ID's).

b) Next, pixel-members that are not re-detected pixels (dark pixels in Fig. 11) are assigned depending on the set of events to which re-detected pixel-members (light grey pixels) of this c.c. have been assigned by rule a) above. Specifically,

i. if all re-detected pixel-members of this c.c. are assigned to the same event, then so are all other pixel-members of this c.c. (i.e. dark pixels of the connected components shown in Fig. 11d-11f are labeled with the same label as the neighboring light-grey pixels);

- ii. otherwise, all pixel-members of this c.c. that are not re-detected pixels (dark pixels) are assigned to new events, one event per pixel (as illustrated in Fig. 11g and 11h).

The GOES-EFD 0.4 Event Tracker algorithm was initially described by Koltunov et al. 2012a.

Figure 11. Illustration of GOES-EFD Event Tracker module rules for assignment of individual fire pixels (light grey and dark grey square patches) to ongoing events ($E1$, $E2$, ...) or new events (E_{new} , E_{new1} , E_{new2} , ...), depending on the pixel spatial connectivity and proximity to recent historic GOES-EFD fire pixels. The text inside each square denotes the event to which this pixel is being assigned by steps 4 and 5 described in section 2.7. Reproduced from Koltunov et al. 2012a.

New events located inside invalid objects flagged by IADC (sect. 2.5) are rejected. Sometimes, due to difficulty in cloud detection, residual misalignment between basis frames and the inspection frame, sensor noise, and for other reasons, a large number of false positive new events may appear in a single inspection frame. Therefore, based on empirical experimentation and given high temporal frequency of input GOES imagery, GOES-EFD 0.4 sets an upper bound for the number of new events in a single frame (one per 10,000 land pixels per 15 min)

and rejects all new events when the bound is exceeded. In general, the upper bound should reflect prior information about the expected frequency of wildfire ignitions and history of detections in the scene. More optimal implementation of this idea requires further research.

3. Wildfire detection timeliness and reliability assessment using California 2006 fire season data

3.1 Study Scene 1 in California and GOES image data

Our test scene that we called Scene 1 earlier in the paper includes central and southern California and western Nevada and occupies nearly 700,000 km² (Fig. 12). This scene includes a range of ecosystems: from semiarid shrublands, conifer dominated forests, annual grasslands, and intensive agriculture to wetland ecosystems. This diversity leads to different fire regimes and dynamics. Furthermore, our test scene is a densely populated and well-monitored region, generally making it very challenging for satellite detection to come before conventional fire identification.

Figure 12. Test Scene 1 in central California used for initial experiments (see text in section 3 for details)

The GOES image time series available for this test spans the period from June to October 2006. The Training mode (sect. 2.2) of the GOES-EFD algorithm used 640 frames acquired in June-August 2. The static basis images included 29 basis images with less than 5% cloud cover from July 15 - August 2 and represented the 24-hour diurnal cycle at 48-min step, on average. Cloud cover in the selected basis images was mapped in a fully automated way by the RCD algorithm, without the optional interactive fine-tuning allowed in the Training mode. The Detection stage of operation was applied to 2852 frames acquired during August 3 - October 1 at variable time steps (~20 min on average). The Detection stage image time series had two significant temporal gaps due to failure of data ingest infrastructure: 8/23 10:30 through 8/28 13:00 Pacific Standard Time (PST); and 9/08 20:30 through 9/25 14:00 PST. The period August 3 – October 1, excluding the data gaps, will be referred to as the Detection period.

The initial validation experiments and analyses described in this section were aimed at the following objectives:

1. Assess detection timeliness by GOES-EFD, with respect to conventional reporting mechanisms and determine baseline false positive rates. The timeliness and reliability of fire detection were evaluated using the following metrics of performance:
 - a. relative detection latency of incidents with respect to their initial report times from conventional sources,
 - b. estimated lower bound on the number of correctly detected wildfire incidents, and
 - c. estimated lower and upper bounds on the number of events that are not wildfire incidents.
2. Compare GOES-EFD and temporally filtered WF-ABBA with respect to detection timeliness and reliability under identical conditions of application.

A meaningful comparison of an optimized operational algorithm WF-ABBA with an initial prototype version of GOES-EFD is not an entirely trivial task not only due to different stages of their development and intended users, but also due to a significant mismatch between the conditions under which these methods were applied. Particularly, during the test period, WF-ABBA was only applied to GOES imagery acquired every 30 min, which contributes to delays in detection and reduces false positives, while GOES-EFD is designed to process all available frames. Therefore to accomplish both objectives, we applied GOES-EFD algorithm at the Detection stage in two modes:

- “GOES-EFD *Regular*”, which is a normal operation mode that processed all 2852 available GOES inspection images; and
- “GOES-EFD *30 min*” that processed images at half-hourly step as WF-ABBA (1494 images) and was thus suitable for comparing results across algorithms.

Furthermore, to facilitate a comparison with respect to incidents (not pixels), WF-ABBA fire pixels were grouped into events using the GOES-EFD Event Tracking algorithm described (sect. 2.7) with the control parameters that previously resulted in the best estimated performance of WF-ABBA (Koltunov et al. 2012a). Finally, we ignored low confidence detections of each algorithm. These included WF-ABBA with a confidence flag of 5 (“low possibility fire”) and detections of GOES-EFD with anomaly confidence $\gamma < 3.5$ (as defined by eq. (6)).

3.2 Validation methodology and datasets

As we mentioned in section 1.1, a methodology that is appropriate for validation with respect to incident detection timeliness and reliability is substantially different from the approach commonly used for active fire product validation with coincident higher-resolution imagery (e.g. Schroeder et al. 2008a, 2008b). This methodology combines geospatial fire records from operational agencies and multitemporal Landsat image analysis to identify new burn scars. The methodology and the datasets we used are discussed in detail elsewhere (Koltunov et al. 2012a) and briefly summarized below.

We have preprocessed, cross-checked, and merged two wildfire geodatabases for year 2006: a true fire perimeter polygon database compiled by the California Department of Forestry and Fire Protection (CAL FIRE) and point and polygon databases created by the Geospatial Multi-Agency Coordination (GeoMAC) group. Both databases are based on the incident report data

from multiple U.S. state and federal agencies, including CAL FIRE, USDA Forest Service, Bureau of Land Management, National Park Service, and other federal and local agencies and departments. The fire incident records in the databases included, in particular, fire final size, initial report and final containment dates, and for some incidents – the initial report hour. After excluding 18 co-occurring (i.e. overlapping in space and time) incidents to avoid optimistic bias in detection timeliness assessment (see Koltunov et al. 2012a), the test sample included 25 fires that started during the Detection period and burned more than 5 acres (2 hectares) over the entire fire lifetime. The initial report hour was recorded for 13 of these 25 incidents.

A fire incident was considered active at time t , if it was reported earlier than $t+3$ hours and was contained later than $t-48$ hours. In this way, we attempted to account for possible delays in fire initial reporting and also for the possibility of residual burning after the fire was deemed contained. Fires with unavailable initial report hour information were assumed active starting 12 am the day they were reported.

Denote p_k the bounding box of the perimeter polygon for a recorded true fire f_k . To match pixels and events to true fires, we defined the spatio-temporal correspondence rules in as follows:

- A detected fire pixel candidate (x, t) centered at a spatial location x detected in a frame acquired at detection time t matches fire f_k , if and only if f_k is active at time t , and at least one of the following conditions holds :
 - x is within a buffer distance b_1 from the final fire perimeter polygon p_k ;

895 ○ x is within a buffer distance b_2 from the final fire perimeter polygon p_k , and
 896 fire f_k is the closest active fire to the fire pixel (x, t) .

897 • An event E matches fire f_k , if at least one pixel assigned to E matches f_k ,

898 where $b_1 = 5.6$ km, which is an average linear size of a GOES 3.9- μm band pixel in north-south
 899 direction in the study area; and $b_2 = 11.2$ km. This two-buffer rule responds to the uncertainty
 900 or limited accuracy of the information about fire actual start and end times. Because of this
 901 uncertainty, when a detected fire pixel coincides in time and is spatially close to more than one
 902 actual incident, it is difficult to determine which one has been detected. In this case, using the
 903 above two-buffer matching rule in addition to excluding co-occurring incidents from the sample
 904 prevents mistaking detection of an older active fire for early detection of a newer ignition.

905

906 The timeliness of detection for a given fire was evaluated using detection latency that is offset
 907 by the time of initial report from conventional sources. The relative latency is defined as
 908 difference between the times of the first alarm by GOES-EFD and the recorded initial report
 909 time. The latency did not account for GOES data delivery (under a minute) and processing time
 910 (under three minutes). All 25 available fires were used to estimate the fraction of eventually
 911 detected incidents. To estimate detection latency down to hours and minutes we used 13
 912 wildfires with recorded report hour. Other fires that were active during the detection period but
 913 started before that period were used to avoid mislabeling pixels as false positives. These
 914 positives were considered neither useful nor harmful with respect to the objectives of early fire
 915 detection (Table 1) and therefore they did not count toward EFD performance metrics.

916

Some of the unmatched and therefore deemed false positive events may actually represent wildfires omitted in the incident report geodatabase we used. Thus, the fraction of unmatched events represents an upper bound on the false alarms. To estimate a lower bound on false alarms, the unmatched events were additionally cross-checked against new burn scars in Landsat imagery. The methods and datasets used to detect new burns are further described in Koltunov et al. 2012a. The unmatched events for which no new burns were found were considered confirmed false positives. However, because the geodatabases we used include only *wildland fires*, these confirmed false positives may include agricultural, residential, and industrial fires of significant intensity. If a new wildfire burn scar appeared in Landsat image following the unmatched geostationary detection, then the event was considered as a possible wildfire omission in the database, i.e. a possible true positive. Subtracting the estimated number of possible true positives from the number of unmatched events provides a lower bound for the number of false alarms (Koltunov et al. 2012a).

3.3 Results

Figure 13 shows the empirical cumulative distribution functions (c.d.f.) of detection latency for GOES-EFD Regular, GOES-EFD 30min, and the temporally filtered WF-ABBA. By definition, the c.d.f. value for a relative latency $\Delta\tau$ represents the percent of fires for which the time difference between the first satellite detection and the recorded initial report is less than $\Delta\tau$ minutes.

Also, for presentation purposes, we defined four non-exclusive groups of detected fire incidents, based on detection timeliness:

1. Eventually detected fires.

2. Detected in less than 1 hour after the initial report.

3. Detected in less than 30 min after the initial report.

4. Detected before the initial report.

The statistics for these groups are presented in Table 2, which also includes interval estimates for false positive new incidents for each algorithm.

Table 2: Summary of performance with respect to new ignition detection for two variants of GOES-EFD version 0.4 and temporally filtered WF-ABBA version 6.1 for 25 California Scene 1 (sect.3, Fig. 12) wildfires occurred during August 3 – October 1, 2006. The timeliness statistics are based on the incident report times compiled from interagency wildfire records provided by California Department of Forestry and Fire Protection (CAL FIRE) and US Geological Survey (USGS) GeoMac program. The false positive intervals in the bottom row were derived by combining wildfire records and bi-temporal Landsat image analysis (sect. 3.2). The false positive events (last row) may include agricultural, residential, and industrial fires.

Detected wildfire incidents	EFD regular	EFD 30 min	WF-ABBA 30 min
eventually	15/25	13/25	11/25
in < 1 hour	10/13	9/13	7/13
in < 30 min	8/13	5/13	2/13
before reported	4/13	3/13	2/13
Total latency reduction for 13 wildfires with recorded report hour	142 min	75 min	45 min
False (non-wildfire) new events	38 to 61	28 to 41	39 to 55

As seen from these data, in the GOES-EFD Regular mode, in which all available GOES images were processed, the algorithm eventually detected 60% of 25 fires. Nearly 77% of the 13 fires with recorded report hour were detected within the first hour, 61% in 30 min, and 31% before they were reported by the conventional sources, with the total of 142 minute reduction of fire incident latency for these incidents (11 minutes, on average). Based on the wildfire records and Google Earth images, the group of 10 undetected fires consisted of 5 grass/pasture fires, 3 understory fires, 1 overstory fire in a sparse forest-shrub mixture on a steep north-facing slope, and 1 undetermined. Five of these omitted fires were very small incidents that burned under 6 ha and five were contained the day they started.

Furthermore, a comparison with WF-ABBA using Figure 13 and Table 2 demonstrates the potential for GOES-EFD to improve timeliness of the initial detection with respect to existing capabilities. Indeed, when both GOES-EFD and WF-ABBA were applied at 30-min step, the number of new ignition commission errors (bottom row of Table 2) by GOES-EFD was about 25% fewer, on average. In addition, GOES-EFD tends to initially identify wildfires earlier, as apparent from the difference in the corresponding latency c.d.f. curves in Figure 13, and from comparing the corresponding rows in Table 2. For example, GOES-EFD detects 28% more incidents within the first hour and more than twice as many – within first 30 minutes. When the false positive new ignitions are nearly the same (cf. column 2 and 4 in Table 2), GOES-EFD provides more than triple reduction of total latency time. The 12 fires missed by GOES-EFD when run at 30 min step were also missed by WF-ABBA; 11 fires were detected by both algorithms, 2 by GOES-EFD only (these fires were contained within 24h), and 0 solely by WF-ABBA.

Figure 13. A comparison of wildfire detection timeliness for two variants of GOES-EFD version 0.4 and temporally filtered WF-ABBA version 6.1 using a cumulative distribution function (c.d.f.) of detection latency relative to documented times of initial reports from conventional sources for 13 California wildfires occurred during August 3 – October 1, 2006. Incident report times were compiled from interagency wildfire data provided by California Department of Forestry and Fire Protection (CAL FIRE) and US Geological Survey (USGS) GeoMac program.

985

986 **4. Discussion and concluding remarks**

987 First responders and fire managers have long requested a low-cost and reliable operational
988 capability to detect wildfire ignitions soon after start. Although GOES satellites have been
989 operationally monitoring active fires for over a decade and provided useful input to the wildfire
990 management community, hot spot detections from GOES have not significantly contributed to
991 reducing wildfire latency before the initial identification in North America, except in remote
992 areas. This could lead one to conclude that GOES, with its large pixel footprint, is not useful
993 for early warning. Yet the capabilities of satellites and sensors are not the only determining
994 factors for detecting ignitions — algorithms are also essential. The only NOAA operational
995 wildfire algorithm for geostationary satellites, WF-ABBA, was designed and optimized for the
996 task of monitoring and characterizing active fires. This type of wildfire remote sensing is most
997 appropriate for user questions and applications for which timeliness of initial detection is a low
998 priority. Earliest possible identification of new ignitions represents a different type of wildfire
999 remote sensing. It has different target objects and performance metrics, and thus can be better
1000 accomplished by a specialized algorithm focusing on subtle anomalies and first signs of change
1001 from “no-fire” to “fire” class.

1002

1003 Toward this goal, we have presented a prototype version of the GOES Early Fire Detection
1004 system algorithm (Version 0.4), which is a part of our ongoing effort toward developing a low-
1005 cost and reliable operational capability for systematic rapid detection and initial confirmation of
1006 new wildfire ignitions. The GOES-EFD algorithm is focusing on detecting subtle changes,
1007 using recently developed multitemporal and contextual tests. The multitemporal tests are based

on the Dynamic Detection Model approach that compensates the observed background signal for changes caused by dynamic environment. This approach is truly generic and sensor-independent (Koltunov et al. 2009). For example, in addition to the current GOES, it has been effective at detecting thermal anomalies in MODIS data (Koltunov & Ustin 2007) and subtle ecosystem disturbances in Landsat image time series with the Ecosystem Disturbance and Recovery Tracker (eDaRT, Koltunov & Ramirez 2012, 2014; Koltunov et al. 2015). Thus, we expect the method to be highly portable and, with modest modifications, applicable not only to the upcoming GOES-R ABI imagery, but also to other current and future Earth-observing satellite sensors.

Algorithms developed for near real-time operations should not be prohibitively expensive computationally. The current sub-optimal, research-mode, and serial MATLAB code implementing GOES-EFD processed Scene 1 in Fig. 12 at an average rate of 2.5 min per frame on a modest desktop computer (2.9GHz dual-core CPU; 8GB RAM). With the proposed system designed to process large areas by independent parts (scenes), and given the growing operational availability of governmental and private distributed computing infrastructures, we project that available computational power will not be a major operational limitation but leave room for developing significantly more computer-intensive versions of the algorithm. Furthermore, because GOES-EFD is not focused or required to provide consistent global coverage, like WF-ABBA, it can be deployed only where operationally necessary, e.g. where societal risks associated with unnoticed ignitions are significant.

1030 Although more research is needed to comprehensively quantify the differences in detection
1031 between the algorithms, in our initial experiments GOES-EFD detected new ignitions
1032 significantly earlier than WF-ABBA and with fewer false alarms. This indicates that GOES-
1033 EFD has the potential to enhance the geostationary information about ignitions and complement
1034 WF-ABBA monitoring capabilities during initial phase of burning. Much of this improvement
1035 stems from GOES-EFD using the scene temporal change analysis as the core approach.

1036 Although in our experiment WF-ABBA did not detect any fires that were not detected by
1037 GOES-EFD, there are situations in which WF-ABBA may be able to detect fires missed by our
1038 algorithm. For example, ignitions followed by undetected clouds can result in a lack of
1039 detections in two consecutive images and thus will be rejected by GOES-EFD temporal filter
1040 that currently requires detection in two consecutive images; however this is not a problem for
1041 WF-ABBA that uses a different definition of temporal consistency in the filtered product.

1042 Also, we anticipate that in highly cloudy conditions the multitemporal detectors relying on
1043 registration, which in turn relies on cloud detection, may not be effective, and therefore a more
1044 stable contextual technique can result in earlier detection. To have a better insight into relative
1045 strengths and weaknesses of the two geostationary wildfire algorithms, new experiments need to
1046 be performed in the future, including more advanced versions of GOES-EFD and more
1047 advanced versions of WF-ABBA, such as version 6.5.

1048

1049 Our first tests of the GOES-EFD have provided experimental evidence that even in the era of
1050 mobile communication and in densely populated areas of the U.S., such as the State of
1051 California, geostationary surveillance could still be a valuable early warning tool – despite the
1052 low spatial resolution of GOES Imager and other factors complicating satellite detection.

Besides occasionally providing the first alarm, thus reducing the risk of out-of-control fire growth, most incidents can be detected by the GOES-EFD system within the first half-hour, which could be a valuable contribution to tactical situational awareness and operation decision making, leading to more optimal resource allocation. These results also indicate that the earliest alarms by GOES-EFD could become routine in less monitored areas. Omitted fires in our experiment were predominantly short-lived grassland or small-scale understory fires that burned up to 5 ha.

Presently, GOES-EFD is a conceptual research prototype algorithm that uses only two thermal channels and should be viewed as a starting point for a more systematic and greater-scale development, optimization, and validation work. Nearly every module of the algorithm has the need and plenty of room for improvement, as we often noted above in the sections describing specific modules. Below, we briefly mention some additional research themes and potential limitations that need to be addressed:

- Optimal partition of a GOES full scan area into GOES-EFD scenes.
- Detection sensitivity in radiance space versus brightness temperature space.
- The GOES-EFD design in which the Training and Detection stages are separated in time allows for these stages to have different degrees of automation. The Detection stage algorithms are fully automatic, which is the only option compatible with real-time operation. In contrast, the Training stage for a given scene is performed relatively rarely (e.g. once a year), thus making it feasible to conduct brief interactive quality control (QC) procedures as an acceptable trade-off between detection performance and complete automation of training. In the presented experiment, one person invested several hours

reviewing the static basis image candidates and replaced a few, based on cloud detection and image registration quality. Examples of other quality vs. automation choices include: review and correct errors in cloud detection or registration for the chosen basis frames; or use a fully automatic algorithm to select optimal basis images (Koltunov & Ustin 2007), thus potentially simplifying QC; or both.

- The GOES-EFD algorithm is not aware of potential active fires in the static basis images, and therefore the algorithm performance in these pixels can be low. A way to address this issue is to use existing operational active fire products from MODIS, VIIRS, and or GOES to mask out confidently detected past fires.
- Regions with frequent high cloud cover represent a significant challenge to GOES-EFD. For example, to start processing an inspection image GOES-EFD needs the entire scene to be available, which can be a limitation in high cloud cover conditions where a pixelwise method, like WF-ABBA is fully applicable. Also in these conditions, the most effective, multitemporal cloud detection method loses accuracy directly, i.e. due to an insufficient number of clear pixels for background estimation, and indirectly, due to less accurate image registration.
- Currently, fire detection near water bodies is limited by combined effect of the more aggressive cloud detection over water pixels and more stringent requirements from a near-cloud anomaly to be flagged as fire.
- While the false positive rates were evaluated in a relatively large sample in our validation experiment (40 days, 2852 images), only a limited number of incidents started during the test period and were suitable for assessing detection timeliness (13 fires) and true positive rates (25 fires). Hence, many factors influencing the system performance

cannot be evaluated by this paper. It is reasonable to assume that system performance will vary depending on geographical area and season. Obviously, the algorithm needs to be tested over larger areas and more extended periods to help better understand these factors and guide its further development.

The central focus of ongoing improvements is placed on reducing false positives. As we discussed in section 1.1, the EFD target objects (ignitions) are extremely rare, compared to the number of fire pixels – the target objects for an active fire monitoring system. Therefore, a method that is reliable at the GOES pixel level (e.g. WF-ABBA) would dramatically loose relative accuracy if applied to early detect ignitions. Although improving relative to WF-ABBA, our false positive rates for new ignitions are currently far from operationally acceptable rates, especially if these detections were to trigger critical management decisions automatically, without independent verification and combination with other available situational evidence. Currently, the major known types of false positives are due to undetected cloud pixels, imperfect image registration, especially in high cloud cover conditions, systematic noise in GOES Images (scan line/cycle shift), among other reasons. Many of these false detections have distinctive spatial-temporal characteristics in GOES imagery, and thus potentially, they can be identified by more advanced versions of the algorithm.

Although we believe algorithm improvement is necessary and attainable, we also recognize that with subtle anomalies sought in real time, omission and commission errors will continue to be an issue, as with any other satellite remote sensing product at low spatial resolution and a very low cost. Therefore, we emphasize again that alerts from GOES should not be reasonably

expected to reach the reliability levels that justify immediately dispatching a standard initial attack team (engines, dozers, crew, airborne assets, etc.). Instead, the roadmap toward operational capacity should include development of appropriate protocols for operational response to the EFD alerts, including verification, initiating an initial attack, elevating the response level, planning resource allocation, acquiring or awaiting for additional information, etc. These protocols should be developed in close coordination with end user partners and take into account their decision making process and spatial/temporal distributions of the risks associated with the potential actions and inaction. To this end, it is also important to continue and expand retrospective validation of the algorithm across different seasons and geographic regions, so that the product quality assurance information will support its informed and sustained use. The combination of satellite alarms, regular retrospective evaluations, and flexible standard response protocols will strengthen the proactive, GOES-aided approach to fire identification over large territories.

4.1 Looking forward to the GOES-R era

In late 2016, a new generation of NOAA geostationary satellites, GOES-R, is scheduled to be launched, presenting a greater opportunity for GOES-EFD to contribute to wildfire disaster reduction. Indeed, compared to the current GOES, the GOES-R Advanced Baseline Imager will provide a three-fold increase in scan speed, four times higher spatial resolution, additional spectral bands, and more accurate pixel geolocation. With these improvements, wildfires will produce much stronger and more frequent signals, thus boosting detection timeliness. In addition, some types of false alarms will be easier to eliminate using the higher temporal resolution, as long as the range of temporal autocorrelation of the false alarm source signal at a given pixel is shorter than the time interval between images. For example, it is easy to show

that the probability of a false alarm due to temporally independent noise reduces exponentially with frame frequency. However, persistent or “slow-moving” sources of false detections in GOES (e.g. reflective cloud edges) will tend to be even more persistent at a GOES-R ABI pixel level. Similarly, in our experiment with current GOES, higher false positive rates were observed when frequency of observation increased. Thus, early fire detection from GOES-R will likely require some algorithm adjustment and optimization.

Another potential difficulty with early fire detection from GOES-R is that, as we mentioned in sect. 2.3.1, the ABI images will be distributed to users in real time only after data remapping into a fixed coordinate system. Because the remapping process involves spatial interpolation, the fire signal can sometimes be reduced and smeared across neighboring pixels, thus effectively increasing their footprint size. Therefore, if image navigation or registration for a specific scene needs to be additionally refined (as can be projected onward based on the GOES I-M and N-P series operation history), this will result in a second interpolation smearing the signal even further. If the original Level-0 data and the image navigation information were available too, all these losses of information could be avoided, and the timeliness and accuracy of geostationary wildfire detection would significantly increase. Unfortunately, at the start of satellite operations, and potentially forever after, data processing algorithms that benefit society by tracking subpixel anomalies, such as WF-ABBA and GOES-EFD, will not be provided real-time access to the original, non-interpolated, Level-0 samples.

Even though neither GOES nor GOES-R is optimal for early wildfire detection, there is no near-future alternative to relying on the NOAA operational geostationary satellites for early wildfire

warning in the Western hemisphere. This is because spatial and temporal resolutions of sensors are not the only factors determining the operational value of the satellite information technology. The operational long-term continuity and reliability are also critical. GOES and GOES-R series missions do provide these advantages, owing to the NOAA open-data policy, decades of excellence in science and engineering support through the agency Cooperative Institutes and partners, and last but not least, the global importance of these missions for applications other than wildfire response, e.g. weather monitoring and other critical mandates. All of that reduces cost of operational wildfire detection. These advantages, however, can be readily complemented by narrower-focus, incidental, or specially tasked missions. Looking forward, we envisage early wildfire warning routinely provided using data from a constellation of satellites that have different combinations of sensitivity, imaging frequency, detection reliability, long-term and short term data availability, and operational redundancy. This multi-satellite data source ensemble will augment other means of wildfire identification, including traditional ground and airborne observations.

4.2 Conclusion

Every year catastrophic wildfires in the United States make news headlines by taking lives, endangering public safety and health, consuming natural resources, and incurring increasingly unsustainable societal costs, only a small fraction of which are the billions of dollars annually spent on suppression. Not every incident can be detected early from GOES. Not every early detected incident can be contained. Yet it is also beyond argument, that every minute of lead time provided by earlier detection reduces the risks due to latent/unconfirmed ignitions and untimely/uninformed decision making by the first responders and other operational authorities.

Over time, these risks tend to translate into losses that are potentially preventable by GOES-EFD, a low-cost systematic fire discovery tool, the initial development of which was discussed in this paper.

These points and initial results we presented argue for a continuing effort to further develop and comprehensively test the GOES Early Fire Detection system within the operational framework of fire management agencies. The resulting regionally optimized ignition products from GOES-EFD will complement the existing fire identification methods and the global monitoring-characterization capabilities provided by the suite of active fire remote sensing algorithms from different platforms.

Acknowledgments

This work was supported by USDA Forest Service, as the main sponsor, and University of California Davis under Cost Share Agreement 10-IA-11130400-009 “Evaluating operational potential of geostationary early fire detection capabilities at regional level”, with contribution by the U.S. Dept. of Homeland Security (Project No. RSID-11-00096; Contract No. HSHQDC-11-C-00158). We thank Ahmad Hakim-Elahi (UC Davis, Sponsored Program Office) and Bruce A. Davis (Dept. Homeland Security, Science & Technology Directorate) for their dedication to support GOES-EFD development. The GOES Imager data readout capability and technical assistance was provided by the CIMIS (California Irrigation Management Information System) program. We also thank Quinn Hart (UC Davis) for his assistance and advice that helped us preprocess GOES GVAR data; Mark Rosenberg (California Department of Forestry and Fire Protection, CAL FIRE) for providing geospatial wildfire data, support, and useful discussions; Mui Lay (UC Davis) for her great help with Landsat image interpretation; George Scheer (UC Davis) for computation and data management support; Mark Finco and Linda R.

Smith (RedCastle Resources, Inc. /USDA Forest Service, RSAC) for assistance with graphics design; Mark Ruminski (NOAA, NESDIS) for helpful discussions of the algorithm. We are grateful to Elaine Prins (UW-Madison, CIMSS) for her highly valuable insights into operational geostationary wildfire detection and aspects of WF-ABBA algorithm, and thank two anonymous reviewers for useful suggestions and constructive criticism that significantly improved the manuscript.

References

- Boeing (2006) GOES-N Data Book, Revision B; available at <http://goes.gsfc.nasa.gov/text/goes.databookn.html>
- Calle, A., Casanova, J. L. and Romo A. (2006) “Fire detection and monitoring using MSG Spinning Enhanced Visible and Infrared Imager (SEVIRI) data”, *Journal of Geophysical Research*, v.111, G04S06.
- Carr, J. L. (2009). Twenty-Five Years of INR, *The Journal of the Astronautical Sciences*, 57(1-2), pp.505–515
- Csiszar, I., Schroeder, W., Giglio, L., Ellicott, E., Vadrevu, K.P. Justice, C. O. and Wind, B. (2014) “Active fires from the Suomi NPP Visible Infrared Imaging Radiometer Suite: Product status and first evaluation results”, *Journal of Geophysical Research-Atmospheres*, 119(2): pp.803-816 .

- 1237 Feltz, J. M., M. Moreau, E. M. Prins, K. McClaid-Cook, and I. F. Brown, (2003): Recent
1238 validation studies of the GOES Wildfire Automated Biomass Burning Algorithm (WF_ABBA)
1239 in North and South America. In *Proceedings of the 2nd International Wildland Fire Ecology
1240 and Fire Management Congress and AMS 5th Symposium on Fire and Forest Meteorology*,
1241 Orlando, Florida, November 16-20, 2003, 6 pp.
1242
- 1243 Giglio, L. (2010). "MODIS Collection5 Active Fire Product User's Guide, Version 2.4".
1244 Available at: http://modis-fire.umd.edu/Documents/MODIS_Fire_Users_Guide_2.4.pdf
1245
- 1246 Giglio, L., Descloitres, J., Justice, C. O., and Kaufman, Y. J., (2003), "An enhanced contextual
1247 fire detection algorithm for MODIS". *Remote Sensing of Environment*, 87, pp.273-282,
1248
- 1249 Grotenhuis, M., Wu, X., Yu, F., Schmit, T., Lindstrom, S., and Cao, C., "On-orbit
1250 characterization of the GOES Imager channel-to-channel co-registration", In Proc. SPIE
1251 v.8510, Earth Observing Systems XVII, 85101T (2012).
1252
- 1253 Hanna, K. 1991. Direct multi-resolution estimation of ego-motion and structure from motion. In
1254 IEEE Workshop on Visual Motion, Princeton, NJ, pp. 156-162.
1255
- 1256 Irani, M. (2002). Multi-frame correspondence estimation using subspace constraints.
1257 *International Journal of Computer Vision*, 48(3), pp.173-194.
1258

- 1259 Justice, C.O., Giglio, L., Roy, D., Boschetti, L., Csiszar, I., Davies, D., Korontzi, S.,
1260 Schroeder, W., O'Neal, K., Morisette, J. (2011) "MODIS-Derived Global Fire Products" in
1261 *Land Remote Sensing and Global Environmental Change*, v.11, Ramachandran, B; Justice, CO;
1262 Abrams, MJ (Eds.); pp. 661-679.
- 1263
- 1264 Koltunov, A. and Ustin S.L. (2007). "Early fire detection using non-linear multitemporal
1265 prediction of thermal imagery", *Remote Sensing of Environment*, v.110(1) pp.18-28.
- 1266
- 1267 Koltunov, A., Ben-Dor, E., Ustin S.L. (2009) "Image construction using multitemporal
1268 observations and Dynamic Detection Models". *International Journal of Remote Sensing*, v.30
1269 (1) pp.57-83.
- 1270
- 1271 Koltunov A. and Ramirez, C. (2012) "Toward near real time automated monitoring of
1272 anomalous sub-annual scale changes in natural ecosystem development", In *ForestSAT 2012*
1273 Conference CD ROM. Sep 11-14 2012, Corvallis, Oregon, USA.
- 1274
- 1275 Koltunov A., Ustin, S. L., Prins, E (2012a) "On timeliness and accuracy of wildfire detection by
1276 the GOES WF-ABBA algorithm over California during the 2006 fire season", *Remote Sensing*
1277 *of Environment*, v. 127, December 2012, pp. 194–209.
- 1278
- 1279 Koltunov, A., Ustin, S.L., Quayle, B., and Schwind, B. (2012b) "GOES Early Fire Detection
1280 (GOES-EFD) System Prototype". In *Proc. ASPRS 2012 Annual Conference*, Sacramento CA,
1281 March 21-23 2012.

1282

1283 Koltunov, A. and Ramirez, C. (2014) “The Ecosystem Disturbance and Recovery Tracker
1284 (eDaRT) system prototype for high-fidelity near-real time ecosystem monitoring”, In:
1285 *ForestSAT 2014*, Trento, Italy, Nov. 2014.

1286

1287 Koltunov, A., Ramirez, C., and Ustin S. L. (2015) “eDaRT: the Ecosystem Disturbance and
1288 Recovery Tracking system prototype supporting ecosystem management in California”, In:
1289 *2015 NASA Carbon Cycle and Ecosystems Joint Science Workshop*, 19-24 April 2015, College
1290 Park, MD.

1291

1292 Martell, D. L. (2015) “A review of recent forest and wildland fire management decision support
1293 systems research”, *Current Forestry Reports*, 1(2), pp. 128-137,

1294

1295 Matson M. and J. Dozier (1981), “Identification of subresolution high temperature sources
1296 using the thermal IR”, *Photogrammetric Engineering and Remote Sensing*, 47, p.1311-1318.

1297

1298 Mazzeo, G., Marchese, F., Filizzola, C., Pergola, N., and Tramutoli, V. (2007), “A multi-
1299 temporal robust satellite technique (RST) for forest fire detection”, *The 2007 International
1300 Workshop on the Analysis of Multi-Temporal Remote Sensing Images*. pp: Z270-5.

1301

1302 NIFC (2014), “Federal Firefighting Costs”, National Interagency Fire Center, Boise, ID, USA;
1303 available at: www.nifc.gov/fireInfo/fireInfo_documents/SuppCosts.pdf

1304

- 1305 NICC (2013) “Wildland Fire Summary and Statistics Annual Report 2013”, National
1306 Interagency Coordination Center (NICC), available at:
1307 www.predictiveservices.nifc.gov/intelligence/2013_Statssumm/annual_report_2013.pdf
1308
- 1309 Prins, E. M., & Menzel, W. P. (1992). Geostationary satellite detection of biomass burning in
1310 South America. *International Journal of Remote Sensing*, 13, pp.2783-2799.
1311
- 1312 Prins, E. M. and Menzel, W. P. (1994). Trends in South American biomass burning detected
1313 with the GOES visible infrared spin scan radiometer atmospheric sounder from 1983 to 1991.
1314 *Journal of Geophysical Research*, 99, 16719–16735.
1315
- 1316 Prins, E. M., Feltz, J.M., Menzel, W.P., & Ward, D.E. (1998). An overview of GOES-8
1317 diurnal fire and smoke results for SCAR-B and 1995 fire season in South America. *Journal of*
1318 *Geophysical Research*, 103 (D24), pp.31.821–31.835.
1319
- 1320 Prins, E., J. Schmetz, L. Flynn, D. Hillger, and J. Feltz. (2001) “An overview of diurnal active
1321 fire monitoring using a suite of international geostationary satellites”. In *Global and Regional*
1322 *Wildfire Monitoring: Current Status and Future Plans* (F. J. Ahern, J. G. Goldammer, C. O.
1323 Justice, Eds.), SPB Academic Publishing, The Hague, Netherlands, pp. 145 - 170.
1324
- 1325 Prins, E. M., Schmidt, C. C., Feltz, J. M., Reid, J. S., Westphal, D. L., and Richardson, K.
1326 (2003). “A two-year analysis of fire activity in the Western Hemisphere as observed with the
1327 GOES Wildfire Automated Biomass Burning Algorithm”. In *12-th Conference on Satellite*

- 1328 *Meteorology and Oceanography*, CD-ROM: Combined Preprints. 83rd AMS Annual Meeting,
1329 Paper P2.28., Long Beach, CA, February 9-13, 2003
1330
- 1331 Prins, E. M., Schmidt, C. C., Brunner, J. C., Hoffman, J. P. Lindstrom, S. S. and J. M. Feltz.
1332 (2010). The global geostationary Wildfire ABBA fire monitoring network. In 17th Conference
1333 on Satellite Meteorology and Oceanography, Annapolis, MD, September 2010, Amer. Meteor.
1334 Soc., P9.13.
1335
- 1336 QFR (2014), *Quadrennial Fire Review*. Prepared by USDA Forest Service
1337 Fire & Aviation Management; U.S. Dept of Interior Office of Wildland Fire; Available online
1338 at http://www.forestsandrangelands.gov/QFR/documents/2014QFR_Brochure_20140313.pdf
1339
- 1340 Schmidt, C. S. and E. Prins. (2003). GOES wildfire applications in the Western Hemisphere.
1341 *Proceedings of the 2nd International Wildland Fire Ecology and Fire Management Congress*
1342 *and AMS 5th Symposium on Fire and Forest Meteorology*, Orlando, Florida, November 16-20,
1343 2003, 4 pp.
1344
- 1345 Schmidt, C. S., J. Hoffman, E. Prins, and S. Lindstrom (2012). GOES-R Advanced Baseline
1346 Imager (ABI) Algorithm Theoretical Basis Document For Fire / Hot Spot Characterization,
1347 Version 2.5, NOAA NESDIS Center for Satellite Applications and Research, July 30, 2012.
1348

- 1349 Schroeder, W., Prins, E. M., Giglio, L., Csiszar, I., Schmidt, C., Morisette, J., and Morton, D.
1350 (2008a). Validation of GOES and MODIS active fire detection products using ASTER and
1351 ETM plus data. *Remote Sensing of Environment*, 112(5), pp.2711–2726.
1352
- 1353 Schroeder, W., Ruminski, M., Csiszar, I., Giglio, L., Prins, E., Schmidt, C., and Morisette, J.
1354 (2008b). Validation analyses of an operational fire monitoring product: The Hazard Mapping
1355 System. *International Journal of Remote Sensing*, 29(20), pp.6059–6066.
1356
- 1357 Schroeder, W., Oliva, P., Giglio, L., Csiszar, I. A. (2014), “The New VIIRS 375 m active fire
1358 detection data product: Algorithm description and initial assessment”, *Remote Sensing of*
1359 *Environment*, 143(5), pp.85–96.
1360
- 1361 Weaver, J. F., Lindsey, D., Bikos, D., Schmidt, C. C., and Prins, E. (2004). Fire detection using
1362 GOES Rapid Scan imagery. *Weather and Forecasting*, pp.19, 496–510.
1363
- 1364 Xu, W., Wooster, M. J., Roberts, G., and Freeborn, P. (2010). New GOES imager algorithms
1365 for cloud and active fire detection and fire radiative power assessment across North, South and
1366 Central America. *Remote Sensing of Environment*, 114(9), 1876–1895.
1367
- 1368 Roberts, G. and Wooster, M. J. (2014). Development of a multi-temporal Kalman filter
1369 approach to geostationary active fire detection & fire radiative power (FRP) estimation. *Remote*
1370 *Sensing of Environment*, 152, pp.392-412.
1371

1372 Verbesselt, J., Hyndman, R., Newnham, G., and Culvenor, D. (2010) Detecting trend and
1373 seasonal changes in satellite image time series. *Remote Sensing of Environment*, 114 (1),
1374 pp.106-115.
1375