Available online at www.sciencedirect.com
SCIENCE DIRECT®
Remote Sensing of Environment 87 (2003) 273–282
Remote Sensing of Environment
www.elsevier.com/locate/rse

An Enhanced Contextual Fire Detection Algorithm for MODIS

Louis Giglio[a,b,*], Jacques Descloitres[a], Christopher O. Justice[c], Yoram J. Kaufman[b]

[a] Science Systems and Applications, Inc., Lanham, MD, USA
[b] NASA Goddard Space Flight Center, Greenbelt, MD, USA
[c] University of Maryland Department of Geography, College Park, MD, USA

Received 27 January 2003; received in revised form 24 May 2003; accepted 31 May 2003

Abstract

Experience with the first 2 years of high quality data from the Moderate Resolution Imaging Spectroradiometer (MODIS) through quality control and validation has suggested several improvements to the original MODIS active fire detection algorithm described by Kaufman, Justice et al. [Journal of Geophysical Research 103 (1998) 32215]. In this paper, we present an improved replacement detection algorithm that offers increased sensitivity to smaller, cooler fires as well as a significantly lower false alarm rate. Performance of both the original and improved algorithm is established using a theoretical simulation and high-resolution Advanced Spaceborne Thermal Emission and Reflection Radiometer (ASTER) scenes. In general, the new algorithm can detect fires roughly half the minimum size that could be detected with the original algorithm while having an overall false alarm rate 10–100 times smaller.

Keywords: Fire; Fire detection; MODIS

1. Introduction

As part of NASA's Earth Observing System (EOS), the Moderate Resolution Imaging Spectroradiometer (MODIS) is carried on both the Terra and Aqua satellites. The MODIS instruments, which began collecting data in February 2000 (Terra) and June 2002 (Aqua), are being used to generate oceanic, atmospheric, and terrestrial data products (Kaufman, Herring, Ranson, & Collatz, 1998; Masuoka, Fleig, Wolfe, & Patt, 1998). Since launch, emphasis has been given to characterizing instrument performance, determining and monitoring the quality of the data products, and undertaking validation (Morisette, Privette, & Justice, 2002). Based on this understanding, improvements have been made to all of the algorithms. The MODIS active fire products fall within the suite of terrestrial products and provide information about actively burning fires, including their location and timing, instantaneous radiative power, and smoldering ratio, presented at a selection of spatial and temporal scales (Justice, Giglio et al., 2002; Kaufman, Justice et al., 1998). A detection algorithm that identifies the active fires within each MODIS swath forms the basis of these products.

Although the original MODIS fire detection algorithm of Kaufman, Justice et al. (1998) functioned reasonably well following several initial postlaunch revisions collectively known as "version 3" (Justice, Giglio et al., 2002), two significant problems limited the overall quality of the product. Firstly, persistent false detections occurred in some deserts and sparsely vegetated land surfaces, particularly in northern Ethiopia, the Middle East, and Central India. Not unexpectedly, most of these were caused by the algorithm's absolute threshold tests. Secondly, relatively small (yet generally obvious) fires were frequently not detected. In response to these problems, we have developed a replacement version 4 contextual algorithm that offers superior sensitivity to smaller, cooler fires and have yielded fewer blatant false alarms. In this paper, we describe this algorithm.

2. Algorithm description

The improved detection algorithm is based on the original MODIS detection algorithm (Kaufman, Justice et al., 1998), heritage algorithms developed for the Advanced Very High Resolution Radiometer (AVHRR) and the Visible and

* Corresponding author. NASA Goddard Space Flight Center, Code 923, Greenbelt, MD 20771, USA. Tel./fax: +1-301-614-6699.
E-mail address: giglio@hades.gsfc.nasa.gov (L. Giglio).

0034-4257/$ - see front matter © 2003 Elsevier Inc. All rights reserved.
doi:10.1016/S0034-4257(03)00184-6

Infrared Scanner (VIRS) (Giglio, Kendall, & Justice, 1999; Giglio, Kendall, & Mack, in press), and experience with the first 2 years of high quality MODIS data.

The algorithm uses brightness temperatures derived from the MODIS 4-and 11-μm channels, denoted by T_4 and T_{11}, respectively. The MODIS instrument has two 4-μm channels, numbered 21 and 22, both of which are used by the detection algorithm. Channel 21 saturates at nearly 500 K; channel 22 saturates at 331 K. Since the low-saturation channel (22) is less noisy and has a smaller quantization error, T_4 is derived from this channel whenever possible. However, when channel 22 saturates or has missing data, it is replaced with the high saturation channel to derive T_4. T_{11} is computed from the 11-μm channel (channel 31), which saturates at approximately 400 K for the Terra MODIS and 340 K for the Aqua MODIS. The 12-μm channel (channel 32) is used for cloud masking; brightness temperatures for this channel are denoted by T_{12}.

The 250-m resolution red and near-infrared channels, aggregated to 1 km, are used to reject false alarms and mask clouds. These reflectances are denoted by $\rho_{0.65}$ and $\rho_{0.86}$, respectively. The 500-m 2.1-μm band, also aggregated to 1 km, is used to reject water-induced false alarms; the reflectance in this channel is denoted by $\rho_{2.1}$. A summary of all MODIS bands used in the algorithm is shown in Table 1.

2.1. Cloud and water masking

Cloud detection was performed using a technique based on that used in the production of the International Geosphere Biosphere Program (IGBP) AVHRR-derived Global Fire Product (Stroppiana, Pinnock, & Grégoire, 2000). Daytime pixels are considered to be cloud-obscured if the following condition is satisfied:

$$(\rho_{0.65} + \rho_{0.86} > 0.9) \text{ or } (T_{12} < 265 \text{ K}) \text{ or}$$
$$(\rho_{0.65} + \rho_{0.86} > 0.7 \text{ and } T_{12} < 285 \text{ K})$$

Nighttime pixels are flagged as cloud if the single condition $T_{12} < 265$ K is satisfied. These simple criteria were found to be adequate for identifying larger, cooler clouds but consistently missed small clouds and cloud edges. One advantage, however, is that fire pixels were never observed to have been mistakenly flagged as cloud. As noted previously (Justice, Giglio et al., 2002; Seielstad, Riddering, Brown, Queen, & Hao, 2002), this problem has been experienced with other cloud masking methods, including the MODIS cloud mask product (Ackerman et al., 1998). Recent improvements in the latter, however, may allow use of the MODIS cloud mask product to be reincorporated into the fire product during a future reprocessing.

Table 1
MODIS channels used in detection algorithm

Channel number	Central wavelength (μm)	Purpose
1	0.65	Sun glint and coastal false alarm rejection; cloud masking.
2	0.86	Bright surface, sun glint, and coastal false alarm rejection; cloud masking.
7	2.1	Sun glint and coastal false alarm rejection.
21	4.0	High-range channel for active fire detection.
22	4.0	Low-range channel for active fire detection.
31	11.0	Active fire detection, cloud masking.
32	12.0	Cloud masking.

Water pixels were identified using the 1-km prelaunch land/sea mask contained in the MODIS geolocation product. Significant errors have been noted in this data set, and an improved water mask is being developed by members of the MODIS Science Team.

2.2. Detection algorithm components

The purpose of the detection algorithm is to identify pixels in which one or more fires are actively burning at the time of the satellite overpass; such pixels are commonly referred to as "fire pixels". As with most other satellite-based fire detection algorithms, our approach exploits the different responses of middle-infrared and long-wave-infrared bands to scenes containing hot subpixel targets (Dozier, 1981; Matson & Dozier, 1981). In particular, the algorithm looks for a significant increase in radiance at 4 μm, in both an absolute sense as well as relative to the observed 11-μm radiance. This characteristic active fire signature is the result of the enormous difference in 4- and 11-μm blackbody radiation emitted at combustion temperatures as described by the Planck function.

The algorithm examines each pixel of the MODIS swath and ultimately assigns it to one of the following classes: *missing data*, *cloud*, *water*, *non-fire*, *fire*, or *unknown*. Pixels lacking valid data are immediately classified as *missing data* and excluded from further consideration. Cloud and water pixels are identified using the previously described cloud and water masks and are assigned to the classes *cloud* and *water*, respectively. The fire detection algorithm considers only those land pixels that remain.

2.2.1. Identification of potential fire pixels

A preliminary classification is used to eliminate obvious non-fire pixels. Those pixels that remain are considered in subsequent tests (described in the next sections) to determine if they do in fact contain an active fire.

A daytime pixel is identified as a *potential fire pixel* if $T_4 > 310$ K, $\Delta T > 10$ K, and $\rho_{0.86} < 0.3$, where $\Delta T = T_4 - T_{11}$. For nighttime pixels, the reflective test is omitted and the T_4 threshold reduced to 305 K. Pixels failing these preliminary tests are immediately classified as *non-fire* pixels.

There are two logical paths through which fire pixels can be identified. The first consists of a simple absolute threshold test. This threshold must be set sufficiently high so that it is triggered only by very unambiguous fire pixels, i.e. those with very little chance of being a false alarm. The

second path consists of a series of contextual tests designed to identify the majority of active fire pixels that are less obvious.

2.2.2. Absolute threshold test

The absolute threshold criterion remains identical to one employed in the original algorithm (Kaufman, Justice et al., 1998):

$$T_4 > 360 \text{ K (320 K at night)} \quad (1)$$

Despite the high daytime threshold, the utility of this test hinges upon adequate sun glint rejection; otherwise, glint-induced false alarms can occur. This issue is addressed in Section 2.2.6. Nighttime pixels are defined as those having a solar zenith angle $\geq 85°$.

2.2.3. Background characterization

In the next phase of the algorithm, which is performed regardless of the outcome of the absolute threshold test, an attempt is made to use the neighboring pixels to estimate the radiometric signal of the potential fire pixel in the absence of fire. Valid neighboring pixels in a window centered on the potential fire pixel are identified and are used to estimate a background value. Within this window, valid pixels are defined as those that (1) contain usable observations, (2) are located on land, (3) are not cloud-contaminated, and (4) are not *background fire* pixels. Background fire pixels are in turn defined as those having T_4>325 K and ΔT>20 K for daytime observations, or T_4>310 K and ΔT>10 K for nighttime observations.

The window starts as a 3×3 pixel square ring around the potential fire pixel. Due to the triangular along-scan response of the MODIS instrument (Kaufman, Justice et al., 1998), the two along-scan pixels adjacent to the potential fire pixel are deemed unreliable and excluded from the background characterization. The ring is increased to a maximum of 21×21 pixels, as necessary, until at least 25% of the pixels within the window have been deemed valid, and the number of valid pixels is at least eight. During this step, an optimized nearest-neighbor search is used to correct for the "bowtie" effect or overlap between MODIS scans (Nishihama et al., 1997). The 21×21 pixel maximum size, though arbitrary, ensures that the background is sampled within ~ 20 km of the potential fire pixel, a scale found empirically to be appropriate for preventing false alarms induced by an unrepresentative selection of background pixels.

The number of valid pixels within the background window is denoted by N_v. During the characterization process, counts are also made and noted of the number of neighboring pixels rejected as background fires (N_f), and the number of neighboring pixels excluded as water (N_w).

If a sufficient number of valid neighboring pixels can be identified, several statistical measures are computed. These are $\bar{T}_4$ and δ_4, the respective mean and mean absolute deviation of T_4 for the valid neighboring pixels; $\bar{T}_{11}$ and δ_{11}, the respective mean and mean absolute deviation of T_{11} for the valid neighboring pixels; and $\overline{\Delta T}$ and $\delta_{\Delta T}$, the respective mean and mean absolute deviation of ΔT for the valid neighboring pixels. The 4-μm brightness temperature mean and mean absolute deviation of those neighboring pixels that were rejected as background fires are also computed and are denoted by $\bar{T}_4'$ and δ_4', respectively. These last two quantities will prove useful for rejecting certain types of false alarms. As suggested by Giglio et al. (1999), we employ the mean absolute deviation as a measure of dispersion rather than the standard deviation since it is more resistant to outliers (Huber, 1981). For contextual fire detection algorithms, this is highly desirable since contamination of the background window by undetected clouds, water, fires, and other sources is not uncommon.

2.2.4. Contextual tests

If the background characterization was successful, a series of contextual threshold tests are used to perform a relative fire detection. These look for the characteristic signature of an active fire in which both the 4-μm brightness temperature (T_4) and the 4- and 11-μm brightness temperature difference (ΔT) depart substantially from that of the non-fire background. Relative thresholds are adjusted based on the natural variability of the background. The tests are:

$$\Delta T > \overline{\Delta T} + 3.5\delta_{\Delta T} \quad (2)$$

$$\Delta T > \overline{\Delta T} + 6 \text{ K} \quad (3)$$

$$T_4 > \bar{T}_4 + 3\delta_4 \quad (4)$$

$$T_{11} > \bar{T}_{11} + \delta_{11} - 4 \text{ K} \quad (5)$$

$$\delta_4' > 5 \text{ K} \quad (6)$$

Of these conditions, the first three isolate fire pixels from the non-fire background. The factor of 3.5 appearing in test (2) is larger than the corresponding factor of 3 in test (4) to help adjust for partial correlation between the 4- and 11-μm observations. Condition (5), which is restricted to daytime pixels, is primarily used to reject small convective cloud pixels that can appear warm at 4 μm (due to reflected sunlight) yet cool in the 11-μm thermal channel. It can also help reduce coastal false alarms that sometimes occur when cooler water pixels are unknowingly included in the background window. Any test based on δ_{11}, however, risks rejecting very large fires since these will increase the 11-μm background variability substantially. For example, over a typical land surface $\delta_{11} \sim 1$ K, whereas for land pixels spanning a large forest fire δ_{11} will routinely exceed 20 K. For this reason, test (6) will be employed to disable test (5) when the background window appears to contain large fires. This situation is recognized by an elevated value of δ_4'; the

presence of background fire pixels (Section 2.2.3) increases this statistic considerably.

2.2.5. Tentative fire detection

We are now in the position to tentatively identify pixels containing active fires. For nighttime fires, this will in fact be an unambiguous, final identification. For daytime pixels, three additional steps are used to help eliminate false alarms caused by sun glint, hot desert surfaces, and coasts or shorelines. These will be described in subsequent sections.

A daytime pixel is tentatively classified as a fire pixel if:

{test (1) is true}

or

{tests (2) − (4) are true and [test (5) or test (6) is true]},

otherwise it is classified as *non-fire*.

A nighttime candidate fire pixel is definitively classified as *fire* if:

{test (1) is true}

or

{tests (2) − (4) are true},

otherwise it is classified as *non-fire*.

For those daytime and nighttime pixels for which the background characterization failed, i.e. an insufficient number of valid neighboring pixels were identified, only test (1) is applied in this step. If not satisfied, the pixel is classified as *unknown*, indicating that the algorithm was not able to unambiguously render a decision.

2.2.6. Sun glint rejection

Sun glint over small bodies of water, wet soil, cirrus cloud, and in rare instances, bare soil can cause false alarms. Sun glint is rejected with a scheme based on that of Giglio et al. (in press), using the angle θ_g between vectors pointing in the surface-to-satellite and specular reflection directions, where

$$\cos\theta_g = \cos\theta_v\cos\theta_s - \sin\theta_v\sin\theta_s\cos\phi. \quad (7)$$

Here, θ_v and θ_s are the view and solar zenith angles, respectively, and ϕ is the relative azimuth angle. A count is made of adjacent water pixels, i.e. the number of water pixels within the eight pixels surrounding the tentative fire pixel, and is denoted by N_{aw}. The following conditions are then evaluated:

$$\theta_g < 2^\circ \quad (8)$$

$$\theta_g < 8^\circ \text{ and } \rho_{0.65} > 0.1 \text{ and } \rho_{0.86} > 0.2 \text{ and } \rho_{2.1} > 0.12 \quad (9)$$

$$\theta_g < 12^\circ \text{ and } (N_{aw} + N_w) > 0 \quad (10)$$

If one or more of these conditions are satisfied, the fire pixel is rejected as sun glint and classified as *non-fire*, otherwise it is classified as *fire*. Condition (8) rejects any fire pixel within the most intense region of glint; detection under this extreme condition is simply too unreliable as the specularly reflected sunlight can elevate T_4 well above 400 K, even over the land surface. Condition (9), which is less strict, looks for the consistently elevated reflectances across multiple bands that are characteristic of sun glint. Condition (10), which is still less strict, rejects fire pixels occurring near water pixels as too risky a prospect in and near regions of sun glint.

2.2.7. Desert boundary rejection

Any surface feature that produces a sharp radiometric transition or edge can potentially cause either errors of omission and commission for any contextual detection algorithm. In the case of the former, a fire located along a boundary may remain undetected since the edge increases the background variability to the point that relative tests incorporating this variability will fail. The latter case can arise when non-fire pixels along the hotter (and/or more reflective) edge of a boundary are incorrectly rejected as background fires during the background characterization phase. This restricts the sample of valid background pixels to those within the cooler (and perhaps darker) side of the boundary, which skews the background statistics toward cooler values (Giglio et al., 1999; Martín, Ceccato, Flasse, & Downey, 1999). In general, it is this second case, namely an error of commission, or "false alarm" along a boundary between hotter and cooler surfaces, that is a much more common problem for contextual fire detection algorithms.

Earlier AVHRR-based contextual algorithms were generally far more susceptible to this problem since they employed thresholds for rejecting background fire pixels that were much lower (e.g. Flasse & Ceccato, 1996; Justice, Kendall, Dowty, & Scholes, 1996). Flasse and Ceccato (1996), for example, rejected background pixels having T_4>311 K and ΔT>8 K, conditions that are frequently satisfied over ordinary land areas; two examples include dry-season African savannas and most deserts. (For simplicity, we have used our 4-μm channel notation in this example, when in fact the corresponding AVHRR channel has a central wavelength of 3.8 μm.) The background-fire rejection thresholds employed in our MODIS algorithm are so high, however, that inadvertent exclusion of non-fire background pixels is almost always restricted to desert areas. For the present MODIS algorithm, therefore, we refer to the problem of eliminating this type of false alarm as *desert boundary rejection*.

To reject false alarms along desert boundaries, one would like the algorithm to identify those cases in which the rejected background fire pixels are ordinary land pixels that happened to satisfy the somewhat arbitrary background-fire rejection thresholds. In this situation, the 4-μm statistics $\bar{T}_4'$ and δ_4' become useful indicators. Over typical daytime desert surfaces, $\bar{T}_4' \approx 335$ K and $\delta_4' \approx 0.5$ K. For a back-

ground containing energetic fire pixels, however, δ_4' will be much larger (40 K and higher is routine), and $\bar{T}_4'$ will be somewhat larger, perhaps 350–380 K. We therefore incorporate several heuristic tests exploiting these trends as a means of rejecting daytime false alarms that can arise along desert boundaries. These are:

$$N_f > 0.1N_v \quad (11)$$

$$N_f \geq 4 \quad (12)$$

$$\rho_{0.86} > 0.15 \quad (13)$$

$$\bar{T}_4' < 345 \text{ K} \quad (14)$$

$$\delta_4' < 3 \text{ K} \quad (15)$$

$$T_4 < \bar{T}_4' + 6\delta_4' \quad (16)$$

If all conditions are satisfied, the fire pixel is rejected as a hot desert boundary surface and classified as *non-fire*, otherwise the pixel undergoes a final coastal false alarm test. Conditions (11) and (12) restrict rejection to cases in which a significant number of background pixels appear to contain background fires, a signature of desert boundary false alarms. Condition (13) simply restricts the remaining tests to bright regions characteristic of deserts. Test (16), which is satisfied only when a tentative fire pixel stands out very strongly against the rejected *background* fire pixels, is not a false alarm rejection test per se. Rather, it permits detection of gas flares, which are frequently located in desert areas that would otherwise be excluded due to the uniformity of the landscape.

2.2.8. Coastal false alarm rejection

Given the contextual nature of the algorithm, it is important to accurately exclude water and mixed water pixels during the background characterization phase. Such pixels are usually cooler than adjacent land pixels during the day. Unknowingly including a sufficient number of water and mixed water pixels in the background window can therefore depress $\bar{T}_4$ and cause a coastal false alarm. Also contributing to this phenomenon is the fact that compared to land, water pixels frequently have lower values of ΔT due to differences in emissivity. Those water and mixed water pixels contaminating the background can therefore decrease $\overline{\Delta T}$ and increase the likelihood that a false alarm will occur.

In many respects, this is merely a special case of the edge problem discussed in the previous section. For the present algorithm, however, such coastal false alarms can occur almost anywhere and are usually not accompanied by the inadvertent rejection of neighboring non-fire pixels described in the previous section. We therefore treat these types of false alarms separately.

As indicated above, the current MODIS land/sea mask contains significant errors in some areas. The bulk of these errors consist of a 1- to 5-km discrepancy along coast and shoreline and small rivers that are missed entirely. In some cases, even much larger water bodies are not masked accurately. An example from eastern Africa is shown in Fig. 1; here, hundreds of pixels are misclassified as land. We have therefore incorporated additional tests to identify cases in which the background window is contaminated with unmasked water pixels. We used a simple test based on the 0.86- and 2.1-μm reflectances and the Normalized Difference Vegetation Index (NDVI) of the valid background pixels, where $\text{NDVI}=(\rho_{0.86}-\rho_{0.65})/(\rho_{0.86}+\rho_{0.65})$. This particular combination was chosen to reduce the likelihood of confusing cloud shadows and burn scars, which also have low reflectances, with water.

Valid background pixels having $\rho_{2.1}<0.05$ and $\rho_{0.86}<0.15$ and an $\text{NDVI}<0$ are considered to be unmasked water pixels, i.e. water pixels incorrectly classified as land in the MODIS land/sea mask. The number of such pixels is denoted as N_{uw}. If test (1) is not satisfied and $N_{uw}>0$, the tentative fire pixel is rejected and classified as *non-fire*, otherwise it is classified as *fire*. This test will periodically reject smaller valid fires, but the much greater reduction in coastal false alarms generally make these errors of omission tolerable on a global basis. Future planned improvements in the MODIS land/sea mask may ultimately render these tests unnecessary.

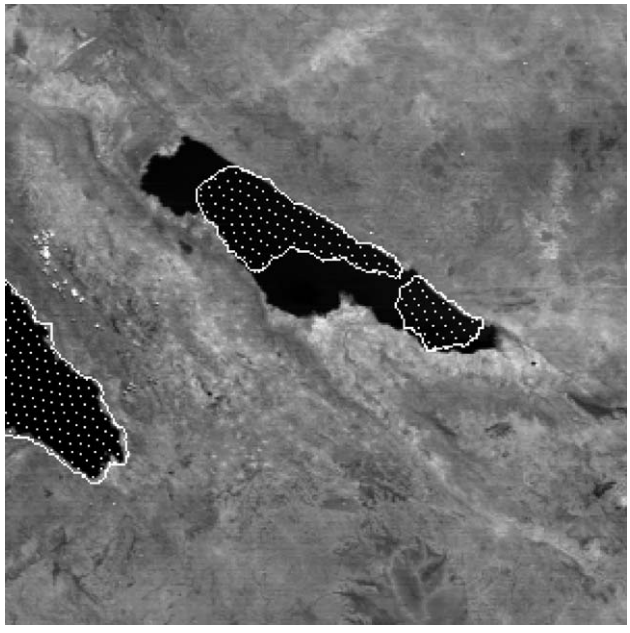

Fig. 1. Example of MODIS water mask (dotted region outlined in white) superimposed on a 1-km aggregated 2.1-μm image of Lake Rukwa in western Tanzania acquired 1 September 2001. The southern tip of Lake Tanganyika appears on the left. Water appears black due to the extreme absorption at this wavelength, and the actual lake boundaries are clearly evident. The image spans an area approximately 300 × 300 km in size.

2.3. Fire detection confidence

A measure of confidence for each detected fire pixel is now also produced based on the approach of Giglio et al. (in press). The measure employs T_4, N_{aw}, the number of cloud pixels adjacent to the fire pixel (N_{ac}), and the standardized variables z_4 and $z_{\Delta T}$, defined as

$$z_4 = \frac{T_4 - \bar{T}_4}{\delta_4} \tag{17}$$

$$z_{\Delta T} = \frac{\Delta T - \overline{\Delta T}}{\delta_{\Delta T}}. \tag{18}$$

These quantities represent the number of absolute deviations that T_4 and T_{11} lie above the background, and are analogous to the more commonly used Z-scores that are calculated using the standard deviation (Devore, 1987). We further employ a ramp function, defined as

$$S(x;\alpha,\beta) = \begin{cases} 0; & x \le \alpha \\ (x-\alpha)/(\beta-\alpha); & \alpha < x < \beta \\ 1; & x \ge \beta \end{cases} \tag{19}$$

The confidence assigned to each fire pixel is composed of a combination of five sub-confidences, labeled C_1 to C_5, each having a range of 0 (lowest confidence) to 1 (highest confidence). For daytime fire pixels, these are defined as

$$C_1 = S(T_4; 310\ \text{K}, 340\ \text{K}) \tag{20}$$

$$C_2 = S(z_4; 2.5, 6) \tag{21}$$

$$C_3 = S(z_{\Delta T}; 3, 6) \tag{22}$$

$$C_4 = 1 - S(N_{ac}; 0, 6) \tag{23}$$

$$C_5 = 1 - S(N_{aw}; 0, 6) \tag{24}$$

For C_1, 310 K represents the minimum brightness temperature required for a pixel to be considered a fire pixel (and is thus less obviously a fire), while based on operational experience, 340 K represents a typical value for a reasonably obvious fire. For C_2, $z_4 = 2.5$ is the minimum value required of fire pixels by the detection algorithm, whereas $z_4 = 6$ represents a typical value (again based on operational experience) for an unambiguous fire. A similar rationale applies to the definition of C_3. C_4 reduces the detection confidence as the number of adjacent cloud pixels increases, accounting for the fact that fire pixels detected along cloud edges are more likely to suffer from cloud contamination, potentially triggering a false alarm via reflected sunlight. Finally, C_5 reduces the confidence as the number of adjacent water pixels increases, reflecting the greater likelihood that the detected fire pixel is instead a coastal false alarm.

Following Giglio et al. (in press), the detection confidence C is then defined as the geometric mean of the sub-confidences, i.e.

$$C = \sqrt[5]{C_1 C_2 C_3 C_4 C_5} \tag{25}$$

For nighttime fire pixels, the thresholds of C_1 are altered appropriately so that

$$C_1 = S(T_4; 305\ \text{K}, 320\ \text{K}), \tag{26}$$

and the cloud- and water-related sub-confidences are not considered. The nighttime detection confidence is therefore simply the geometric mean of C_1, C_2, and C_3.

3. Algorithm performance

To date, four principal methods have been used to assess algorithm performance and evaluate the MODIS fire products. First, based on earlier work done by Dowty (1993) and Giglio et al. (1999), simulated MODIS imagery was used to quantify algorithm detection and false alarm rates under a wide range of environmental conditions within different biomes. Second, fire maps generated from high-resolution scenes acquired with the Advanced Spaceborne Thermal Emission and Reflection Radiometer (ASTER) have been compared to fires identified by versions 3 and 4 algorithms. Third, the U.S. Forest Service Fire Sciences Laboratory, which has undertaken an independent validation of the MODIS fire product, has found good correspondence between MODIS fire locations and fire perimeters measured by the Forest Service. Finally, in a more qualitative approach, unambiguous fire pixels and obvious false alarm sources were identified through visual inspection of MODIS 250-, 500-m, and 1-km imagery, permitting the output from the original and improved algorithms to be compared to "expert"-derived fire masks. Although this approach is in general greatly inferior to the more rigorous approaches mentioned above, it does allow obvious problems to be diagnosed and corrected.

In Section 3.1, we will describe our simulation approach and then discuss the theoretical algorithm performance established by way of the simulation. In Section 3.2, we will briefly report on the recent results of Morisette, Giglio, Csiszar, and Justice (2003), who have completed the first of many planned regional MODIS fire validation activities. (Ultimately, the MODIS fire products will be validated globally through this process.) Finally, in Section 3.3, we will present two examples of obvious cases in which a simple visual analysis and comparison is appropriate.

3.1. Simulated fire scenes

3.1.1. Method

Simulated 25 × 25-km images of MODIS channels 21, 22, and 31 were generated using MODIS-specific modifications of the method used by Giglio et al. (1999) in their evaluation of several AVHRR active fire detection algorithms. The scenes depicted in these images contained idealized fires of various sizes and temperatures in 10 different biomes (desert, tropical rainforest, tropical deciduous forest, tropical savanna, temperate deciduous forest, temperate evergreen forest, temperate grassland, boreal evergreen forest, boreal deciduous forest, and tundra). Each biome was characterized by a range of seasonal average surface temperatures and seasonal average emissivities at 1-km spatial resolution. Following Giglio et al. (1999), the 1-km grid cells of each 25 × 25-km scene were assigned individual 4- and 11-μm emissivities and a temperature by drawing random samples from normal distributions. In this manner, 100 model surfaces were generated for each biome and season. The MODTRAN atmospheric model (Berk, Bernstein, & Robertson, 1989) was used to calculate atmospheric transmission and thermal contributions with an appropriate seasonal tropical, temperate, mid-latitude, or sub-Arctic atmospheric model. A sensor view was then computed taking into account the shape and overlap of the MODIS pixels that fitted inside the 25-km surface grid.

The fire detection algorithm was applied to the synthetic MODIS imagery, and by repeating the process over a range of conditions, the algorithm's performance was characterized statistically in terms of probability of fire detection (P_d) and false alarm (P_f). Both probabilities are functions of fire temperature and area, solar and viewing geometry, visibility, season, and biome. The presence of fires in the background window also affects P_d and P_f (Giglio et al., 1999).

None of the reflective channels (channels 1, 2, and 7) are currently modeled in the simulation, so the tests involving these channels were omitted in the application of the algorithm to the simulated data. Since these tests prevent highly reflective and coastal pixels from being classified as fire, the effect of their omission was to potentially overestimate both P_d and P_f. This approach nevertheless yields useful upper bounds for both probabilities.

3.1.2. Results

Because the probability of detection is so strongly dependent upon the temperature and area of the fire being observed, P_d is summarized as a *detection matrix* in which fire temperature and area form the rows and columns of the matrix. Such matrices are shown graphically in Fig. 2.

Over all biomes considered, the size of the smallest flaming fire having at least a 50% chance of being detected under both ideal daytime and nighttime conditions was ~ 100 m^2. For the version 3 algorithm, this value was more than two times larger. We define ideal conditions to mean that the fire is observed at or near nadir on a fairly homogeneous surface, the background window contains no fires, and the scene is free of clouds, heavy smoke, and sun glint. For nighttime cases in the coldest biomes, this minimum area tended to be somewhat larger, typically by a factor of two, since the "universal" thresholds used to identify potential fire pixels (Section 2.2.1) become much

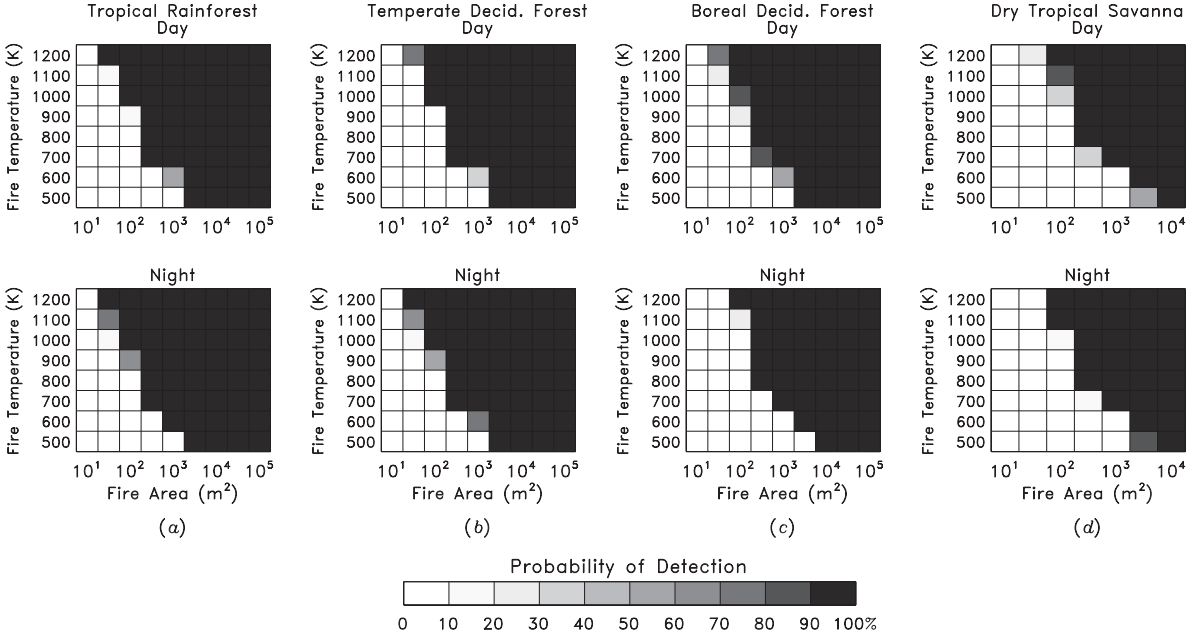

Fig. 2. Nadir (0° scan angle) daytime (top row) and nighttime (bottom row) detection matrices for four different biomes. (a) Tropical rainforest with 0° daytime solar zenith angle; (b) temperate deciduous forest (summer) with 0° daytime solar zenith angle; (c) boreal deciduous forest (summer) with 0° daytime solar zenith angle; (d) dry-season tropical savanna with 40° daytime solar zenith angle.

less appropriate under these conditions. Purely smoldering fires generally had to be 10–20 times larger to achieve a similar probability of detection.

For the improved algorithm, no false detections were observed under any circumstances. This is in contrast to the version 3 algorithm, which produced false alarms in daytime desert scenes at solar zenith angles below 20°. As mentioned previously, this difference is due to the fact that "risky" absolute threshold tests are not employed in the new algorithm.

An interesting case is the daytime dry-season tropical savanna (Fig. 2d). Earlier work by Giglio et al. (1999) and Giglio et al. (in press) showed that contextual algorithms using data from either the AVHRR or the VIRS instruments have almost no capability to detect fires at small or even moderate solar zenith angles, e.g. $\theta_s < 40°$. For both instruments, this is caused by the relatively low saturation (~ 325 K) of their mid-infrared channels; fires show little or no contrast against the hot, bright savanna surface that can saturate the mid-infrared channel even in the absence of a fire. The high saturation of the MODIS band 21, however, allows detection to proceed largely unhampered.

3.2. Validation using ASTER scenes

Recent work by Morisette et al. (in press) used 18 high-resolution ASTER scenes to validate the MODIS fire product in southern Africa. The ASTER instrument is carried on-board the Terra satellite, allowing spatially and temporally coincident observations to be acquired with those of the Terra MODIS. This is ideal for fire validation since active fires can alter a scene significantly in a relatively short period of time (e.g. minutes). Although this first investigation was limited to southern Africa, additional work is ongoing to assess the MODIS fire product globally using a much larger collection of ASTER scenes.

Using the ASTER 2.4-μm channel, which has a spatial resolution of 30 m, Morisette et al. (in press) were able to map the "true" distribution of active fires in each scene. The 30-m pixels of the resulting ASTER fire masks were then assigned to those MODIS pixels in which they fell; this was accomplished by performing a nearest-neighbor search through all MODIS pixels overlapping the region viewed within the particular ASTER scene. In this manner, ASTER fire masks were prepared for a total of 66,761 MODIS pixels. The probability of detection was then related to the absolute number of ASTER fire pixels (a proxy for instantaneous fire size) within a MODIS pixel. The spatial heterogeneity of the ASTER fire pixels was considered as well. The conditional number of ASTER fire pixels observed in a single MODIS pixel varied from 1 to more than 300. Cases in which no ASTER fire pixels were observed in the corresponding MODIS pixel were also included.

For the time period that was considered (5 August–6 October 2001), Morisette et al. found that the minimum fire size to achieve $P_d \geq 0.5$ was, in terms of number of ASTER pixels, approximately 30. In contrast, the original detection algorithm required a minimum of 48 ASTER fire pixels regardless of the statistical model. Using the paired ASTER–MODIS data prepared by Morisette et al. (in press), which were available to us, we performed a subsequent analysis of each algorithm. With respect to false alarms, the algorithms detected 4 (version 3) and 12 (version 4) fire pixels for which the corresponding ASTER fire mask showed zero. We do not, however, consider these to be true false alarms. In each case, ASTER fire pixels were present in at least one adjacent along-scan or along-track MODIS pixel, suggesting that these apparent false alarms might be

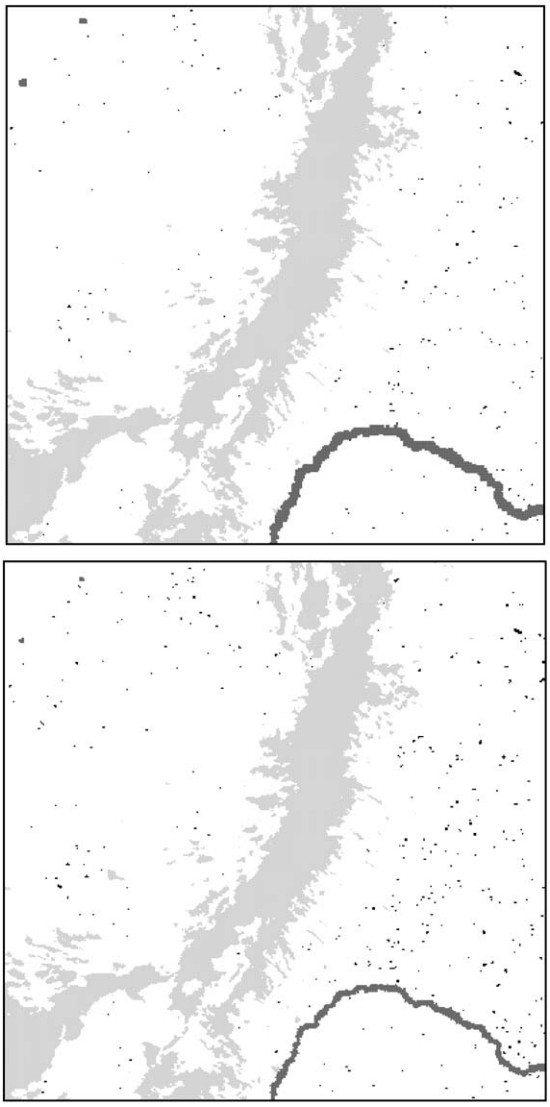

Fig. 3. Example of small fires detected by original (version 3) algorithm (top) and version 4 algorithm described in this paper (bottom) acquired over the northern portion of the Democratic Republic of the Congo at 09:30 on 16 December 2000. Fire pixels are shown in black, rivers in dark gray, clouds in light gray, and non-fire areas in white. The original algorithm detected 267 fire pixels, whereas the new algorithm detected 568. The image spans a region approximately 400 km along each side. Note that for the version 3 algorithm, shorelines were expanded to reduce coastal false alarms; this step is unnecessary for the improved detection algorithm.

caused by limitations in the 2-km rectangular model of the MODIS pixel footprint employed by Morisette et al. In addition, Terra MODIS channel 21 exhibits residual instrument artifacts, including minor "blooming" near large or very hot fires (Justice, Giglio et al., 2002), which could induce false alarms in the immediate vicinity of fires. A better instrument characterization is needed to quantify the extent of these residual artifacts. A third possibility is that the algorithm is simply labeling hot, adjacent, recently burned patches as fires. The corresponding ASTER scenes showed little or no signs of burn scars in these MODIS pixels, however, and we therefore view this explanation as unlikely.

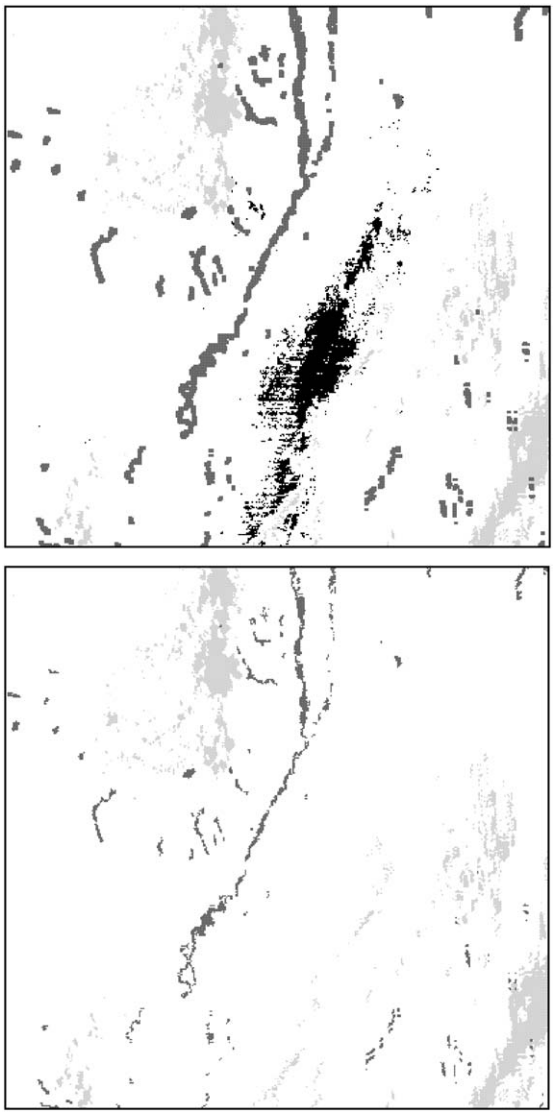

Fig. 4. Example of blatant false fire pixels detected by original (version 3) algorithm (top) in Pakistan on 13 June 2001, 06:30 UTC. Version 4 algorithm output is shown in the bottom figure. Fire pixels are shown in black, rivers in dark gray, clouds in light gray, and non-fire areas in white. Version 3 algorithm detected 4759 false fire pixels; version 4 algorithm detected zero. The image spans a region approximately 400 km along each side.

3.3. Visual comparison

An example of the improved ability of the version 4 algorithm to detect small fires is shown in Fig. 3. In this example from the northern part of the Democratic Republic of the Congo, the version 4 algorithm detected more than twice the number of fire pixels detected with version 3. A detailed visual analysis of the MODIS visible, mid-infrared, and long-wave-infrared bands for this scene indicates that the hundreds of additional fire pixels—most of which are accompanied by distinct smoke plumes—are in fact true fires rather than false alarms.

As mentioned in Section 1, the original detection algorithm suffered from persistent false detections in deserts and other sparsely vegetated land surfaces. An example over Pakistan is shown in Fig. 4. Here, the version 3 algorithm yielded nearly 4800 false fire pixels. These pixels were deemed false based on an examination of MODIS 250-, 500-m, and 1-km imagery and the fact that the large cluster persisted over long periods of time (e.g. weeks). All fire pixels were located in areas of sparsely vegetated soil, and none of the fire pixels have sufficiently high 4-μm brightness temperatures to definitively suggest that true fires might be present. In addition, the top-of-atmosphere visible channel reflectances reveal a complete absence of smoke. This suggests that the majority (if not the entirety) of the fire pixels within this scene are indeed false alarms. The corresponding output from the version 4 algorithm, which yielded no false detections, is shown for comparison. Although the difference in algorithm output is enormous in this particular example, the average reduction in false alarm rate at the global scale is generally not this dramatic. Test runs using almost a year of MODIS data indicated that the number of obvious false fire pixels generated with the version 4 algorithm is generally 10–100 times smaller than the number generated with version 3.

4. Conclusion

We have described an improved contextual active fire detection algorithm for the MODIS instrument. This algorithm, known as version 4, offers considerable improvement over previous versions. The version 4 algorithm is run as part of the MODIS land forward processing stream, as well as within the MODIS Rapid Response System (Justice, Townshend et al., 2002). It is also being run as part of the MODIS "Collection 4" reprocessing stream to reprocess all MODIS data starting from March 2000, the beginning of the Terra MODIS data archive. It should be noted that the Terra instrument performance prior to November 2000 renders the active fire data for this period of limited utility. Consistent time series of fire data from MODIS should be started no earlier than November 2000. The same version 4 algorithm applies to both Terra and Aqua MODIS data.

Although additional algorithm enhancements may be made in the future based on validation results or changes in instrument performance, the immediate focus must now be given to improving the ancillary water mask used in the algorithm. This is currently the greatest source of error in the version 4 product. Commission errors in the current water mask cause some persistent false alarms along the banks of (and islands within) some rivers. Improvements in the land–water mask derived either directly from MODIS, for example, by using the 250-m bands, or from an external data source, such as the emerging data sets from the Shuttle Radar Topography Mission, can be envisioned.

Ongoing efforts are being made to assess the detection algorithm performance and determine product accuracy under different conditions. For the Terra fire product, validation with ASTER is actively being pursued and globally representative validation is underway. Improvements are also being made to the fire simulation to model all bands and add greater realism to the simulation data set. There is also a need to better model some of the idiosyncrasies of the MODIS bands 21 and 22 detectors.

Iterative improvements to the MODIS land algorithms, followed by strategic data reprocessing, will lead to the long-term science quality data products needed for global change research. During the first 2 years of MODIS data, we have gained considerable experience and understanding of the fire algorithm and the Terra MODIS instrument performance, which has resulted in the significant improvements to the algorithm presented in this paper.

Acknowledgements

We thank Jeffrey Morisette (NASA Goddard Space Flight Center) and four anonymous reviewers for their helpful comments on the manuscript.

References

Ackerman, S. A., Strabala, K. I., Menzel, W. P., Frey, R. A., Moeller, C. C., & Gumley, L. E. (1998). Discriminating clear sky from clouds with MODIS. *Journal of Geophysical Research*, *103*, 32141–32157.

Berk, A., Bernstein, L. S., & Robertson, D. C. (1989). MODTRAN: A moderate resolution model for LOWTRAN 7. Final Report GL-TR-89-0122, Geophysics Laboratory, U.S. Air Force Systems Command, Hanscomb AFB, Massachusetts, USA.

Devore, J. L. (1987). Probability and statistics for engineering and the sciences. Monterey: Brooks/Cole Publishing. 672 pp.

Dowty, P. R. (1993). A theoretical study of fire detection using AVHRR data. MS Thesis. Department of Environmental Sciences, University of Virginia, Charlottesville, Virginia, USA.

Dozier, J. (1981). A method for satellite identification of surface temperature fields of subpixel resolution. *Remote Sensing of Environment*, *11*, 221–229.

Flasse, S. P., & Ceccato, P. (1996). A contextual algorithm for AVHRR fire detection. *International Journal of Remote Sensing*, *17*, 419–424.

Giglio, L., Kendall, J. D., & Justice, C. O. (1999). Evaluation of global fire detection algorithms using simulated AVHRR infrared data. *International Journal of Remote Sensing*, *20*, 1947–1985.

Giglio, L., Kendall, J. D., & Mack, R. (2003). A multi-year fire data set for the tropics derived from the TRMM VIRS. *International Journal of Remote Sensing* (in press).

Huber, P. (1981). *Robust statistics*. New York: Wiley.

Justice, C. O., Giglio, L., Korontzi, S., Owens, J., Morisette, J. T., Roy, D., Descloitres, J., Alleaume, S., Petitcolin, F., & Kaufman, Y. (2002). The MODIS fire products. *Remote Sensing of Environment*, *83*, 244–262.

Justice, C. O., Kendall, J. D., Dowty, P. R., & Scholes, R. J. (1996). Satellite remote sensing of fires during the SAFARI campaign using NOAA AVHRR data. *Journal of Geophysical Research*, *101*, 23851–23863.

Justice, C. O., Townshend, J. R. G., Vermote, E. F., Masuoka, E., Wolfe, R. E., Saleous, N., Roy, D. P., & Morisette, J. T. (2002). An overview of MODIS land data processing and product status. *Remote Sensing of Environment*, *83*, 244–262.

Kaufman, Y. J., Herring, D. D., Ranson, K. J., & Collatz, G. J. (1998). Earth Observing System AM1 mission to Earth. *IEEE Transactions on Geoscience and Remote Sensing*, *36*, 1045–1055.

Kaufman, Y. J., Justice, C. O., Flynn, L. P., Kendall, J. D., Prins, E. M., Giglio, L., Ward, D. E., Menzel, W. P., & Setzer, A. W. (1998). Potential global fire monitoring from EOS-MODIS. *Journal of Geophysical Research*, *103*, 32215–32238.

Martín, M. P., Ceccato, P., Flasse, S., & Downey, I. (1999). Fire detection and fire growth monitoring using satellite data. In E. Chuvieco (Ed.), *Remote sensing of large wildfires* (pp. 101–122). Berlin: Springer-Verlag.

Masuoka, E., Fleig, A., Wolfe, R. W., & Patt, F. (1998). Key characteristics of the MODIS data products. *IEEE Transactions on Geoscience and Remote Sensing*, *36*, 1313–1323.

Matson, M., & Dozier, J. (1981). Identification of subresolution high temperature sources using a thermal IR sensor. *Photogrammetric Engineering and Remote Sensing*, *47*, 1311–1318.

Morisette, J. T., Giglio, L., Csiszar, I., & Justice, C. O. (2003). Validation of the MODIS Active fire product over southern Africa with ASTER data. *International Journal of Remote Sensing* (in press).

Morisette, J. T., Privette, J. L., & Justice, C. O. (2002). A framework for the validation of MODIS land products. *Remote Sensing of Environment*, *83*, 77–96.

Nishihama, M., Wolfe, R., Solomon, D., Patt, F., Blanchette, J., Fleig, A., & Masuoka, E. (1997). MODIS level 1A Earth location: Algorithm theoretical basis document version 3.0. SDST-092, MODIS Science Data Support Team.

Seielstad, C. A., Riddering, J. P., Brown, S. R., Queen, L. P., & Hao, W. M. (2002). Testing the sensitivity of a MODIS-like daytime active fire detection model in Alaska using NOAA/AVHRR infrared data. *Photogrammetric Engineering and Remote Sensing*, *68*, 831–838.

Stroppiana, D., Pinnock, S., & Grégoire, J. -M. (2000). The Global Fire Product: Daily fire occurrence from April 1992 to December 1993 derived from NOAA AVHRR data. *International Journal of Remote Sensing*, *21*, 1279–1288.