A Hybrid Contextual Approach to Wildland Fire Detection Using Multispectral Imagery

Ying Li, Anthony Vodacek, Robert L. Kremens, Ambrose Ononye, and Chunqiang Tang, *Member, IEEE*

Abstract—We propose a hybrid contextual fire detection algorithm for airborne and satellite thermal images. The proposed algorithm essentially treats fire pixels as anomalies in images and can be considered a special case of the more general clutter or background suppression problem. It utilizes the local background around a potential fire pixel and discriminates fire pixels based on the squared Mahalanobis distance in multispectral feature space. It also employs the normalized thermal index to identify background fire pixels that should be excluded from the calculation of the statistical properties of the local background. The use of the squared Mahalanobis distance naturally incorporates the covariance of the multispectral image into the decision and requires the setting of a single detection threshold. By contrast, previous contextual algorithms only incorporate the statistical properties of individual bands and require the manual setting of multiple thresholds. Compared with the latest Moderate Resolution Imaging Spectroradiometer fire product (version 4), our algorithm improves user accuracy and producer accuracy by 1.5% and 2.6% on average, respectively, and up to 28% for some images. In addition, the novel use of the squared Mahalanobis distance allows us to create fire probability images that are useful for fire propagation modeling. As an example, we demonstrate this use for the airborne data.

Index Terms—Anomaly detection, Mahalanobis distance, multispectral images, wildland fire detection.

I. INTRODUCTION

THE effects of wildland fire are important at both local scales where impacts to human safety and property become critical and at regional or global scales where wildland fire can have important climate and ecological effects via gas and particle emissions [1]–[4] and changes to plant communities [5], [6]. Because of the dangerous environment posed by wildland fires, active fire distribution and temporal variation is perhaps best assessed on a local scale using airborne imaging and on a global scale using satellite imaging. Further, the continued development of models for forecasting wildland fire behavior and propagation can benefit from the use of high-resolution wildland fire images from an airborne platform as a data source for initializing and nudging model predictions [7], [8].

The most common method for detecting fires relies on the detection of the Planck emission from fire in the midwave infrared

Y. Li, A. Vodacek, R. L. Kremens, and A. Ononye are with the Center for Imaging Science, Rochester Institute of Technology, Rochester, NY 14623 (e-mail: yxl8528@cis.rit.edu).

C. Tang is with the IBM T. J. Watson Research Center, Yorktown Heights, NY 10598 USA.

(MWIR) and the longwave infrared (LWIR) [9]–[12]. Both airborne [13]–[16] and satellite [10], [17], [18] remote sensing systems with the appropriate bands have been used to study wildland fire in the past several decades. Although visual analysis has remained important for operational use [16], automated algorithms have been developed for fire detection on a global scale from satellite images and will need to be developed for real time airborne applications such as fire propagation modeling. Some users are interested in airborne data, since the ground spot diameter (GSD) will be suitable for most operational fire detection work, for which fire locations need to be specified to within tens of meters. To date, the most flexible approach to fire detection is represented by the contextual algorithms such as the one used for the Moderate Resolution Imaging Spectroradiometer (MODIS) fire product [10], [11]. Contextual algorithms consider the local neighborhood of the pixel under examination to cope with the variations in the background environmental temperature that may occur across the large area viewed by satellite systems. For high-resolution airborne data, the range of background temperature due to microclimates can introduce as much variability as small fires. Existing contextual algorithms calculate the statistical properties of a single MWIR band, a single LWIR band, and then use a set of threshold tests of those bands to identify fire pixels. These algorithms have several limitations. First, for a specific sensor, these algorithms need manually tuning of a series of thresholds according to the available spectral bands. Second, these algorithms do not take advantage of the important inter-band information of the multispectral data that may be available for a given sensor.

We propose a hybrid contextual algorithm for fire detection to overcome these limitations. The proposed algorithm is applicable to multispectral near infrared, and thermal infrared data. It still falls into the category of contextual algorithms, but it distinguishes itself with the use of an improved technique for defining the local neighborhood. More importantly, it uses the well-known squared Mahalanobis distance operating in multispectral feature space to identify fire pixels. Implicit in the squared Mahalanobis distance is a degree of directionality in the feature space based on the covariance of the features for the background noise. Our algorithm can be considered as a special case of the more general clutter or background suppression problem [19]. Our algorithm has several major advantages over existing contextual algorithms.

1) It is based on the nonstationary mean nonstationary covariance image model [20], and can be easily applied to different sensors that use different band sets. We demonstrate the flexibility of our algorithm with images taken by two different types of sensors: the MODIS sensors and our

own Wildfire Airborne Sensor Project (WASP) sensor. By contrast, existing algorithms typically only work for specific sensors because they employ many manually constructed rules that only deal with properties of individual bands. As a result, a large number of algorithms have been proposed for different sensors.

2) Our algorithm is relatively simple. It has only two parameters that can be easily tuned and includes no manually constructed rules.

3) The use of the squared Mahalanobis distance in our algorithm has the added advantage of providing a direct means of deriving a fire probability value that is required for new approaches to fire propagation modeling [8].

4) Our feature-based algorithm does not need calibrated radiance as input. This extends the utility of the method to operational sensors that may not have the laboratory capabilities for producing radiance calibrations.

The rest of the paper is organized as follows. We present the theoretical background of our image model in Section II and derive the squared Mahalanobis distance metric for identifying fire pixels in Section III. The details of our algorithm are described in Section V. We compare our algorithm with the MODIS product in Section VI and provide concluding remarks in Section VII.

II. NONSTATIONARY MEAN, NON-STATIONARY COVARIANCE IMAGE MODEL

The theoretical development of the statistical approach relies on the formulation of a statistical model for the background clutter from which the fires are to be discriminated. In this paper, it is assumed that the background image is a nonstationary mean nonstationary covariance Gaussian process [20]. The nonstationary mean describes the gross background structure of the image and the nonstationary covariance characterizes detailed variation of the image which is position sensitive. Other studies have shown the feasibility of this approach, where images have been modeled as consisting of approximately Gaussian intensity fluctuates about a nonstationary mean with a nonstationary covariance, especially for thermal images [21], [22]. A nonstationary mean background can be estimated by blurring the image using a low-pass filter, after masking background target pixels (fire pixels in this case) so they are not included in the calculation. We use the normalized thermal index (NTI) [12] for determining background fire pixels to exclude them prior to application of the low-pass filter (see Section V-B). Since the covariance of the image varies slowly relative to the mean, when the blurred image is subtracted from the original image, the residual image in a local window, where the covariance remains unchanged, can then be modeled as a Gaussian process fluctuating around a mean of zero.

We tested this assumption with the MODIS data collected on June 18, 2002 over the western U.S. (MODIS granule ID: MOD021KM.A2 002169.1740.003.2002170040927.hdf). Fig. 1(a) is a portion of band 20 (3.660–3.840 μm) image of the MODIS granule data mentioned above with the size of 450×400 pixels. Fig. 1(b) is the residual image created by first using a 9×9 low-pass filter (i.e., sliding window average) to
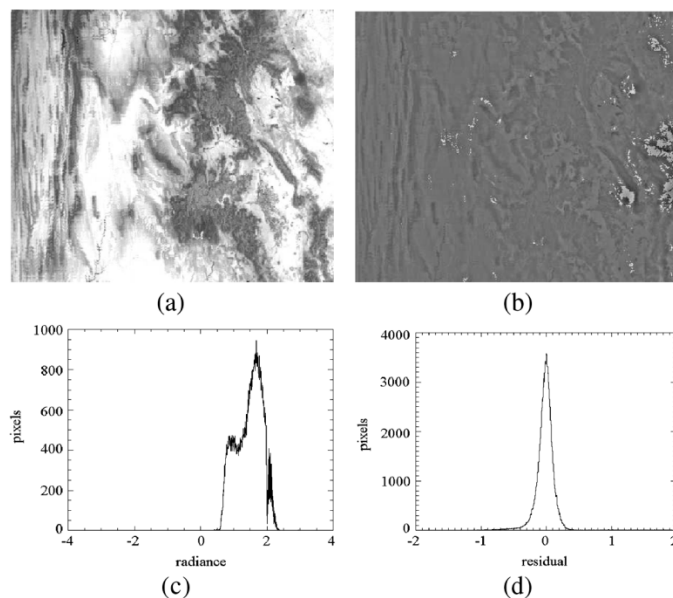

Fig. 1. (a) MODIS band 20 image of Colorado, acquired on June 18, 2002. (b) The residual image of the image in (a) after a 9×9 sliding window mean removal. (c) Histogram of the image in (a). (d) Histogram of the residual image in (b).

generate a blurred version of the original image and then subtracting the blurred image from the original image. It is shown that one method to obtain an optimal kernel size is by minimizing the third moment of the residual image among 3×3, 5×5, 7×7, 9×9, and 11×11 [23]. From our experiments the kernel size has almost no effect on the end result, so it is not critical to decide the kernel size for this step. We chose 9×9 so that the kernel was sufficiently large to avoid situations where a 3×3 kernel, for example, may have too many background fire pixels to accurately determine the background mean. The position sensitive mean of the original image is manifest in the bimodal histogram shown in Fig. 1(c). In contrast, the residual image has a histogram centered on zero, showing that the background was effectively removed through the use of the 9×9 low-pass filter Fig. 1(d). The residual image is still a nonstationary process. The shape of the histogram, however, is much closer to a Gaussian distribution. This image characteristic is the motivation for the development of contextual fire detection algorithms. To ensure that the covariance remains unchanged, the statistics are calculated for a relatively small region around the target pixel. The local background region used to calculate the covariance is 31×31 pixels. Fig. 2(a) presents the histogram of the residual image in a 31×31 pixel window which is known to contain no fire. Fig. 2(b) shows the associated normal probability plot. The histogram and normal probability plot justify the use of Gaussian processes to model images.

III. DERIVING THE SQUARED MAHALANOBIS DISTANCE METRIC FOR FIRE DETECTION

The fire detection problem for multispectral images can be described using a statistical approach as follows. Suppose one is given an image of N independent pixels, each pixel is composed of J features. The J features can be spectral band radiance or
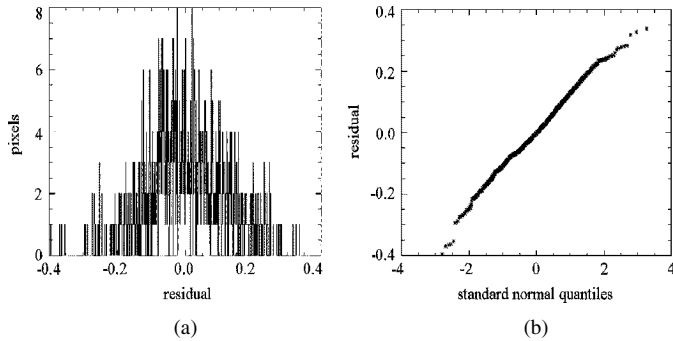
Fig. 2. (a) Histogram of a 31×31 pixel subimage of the residual image in Fig. 1(b). (b) Normal probability plot of the 31×31 pixel subimage of the residual image in Fig. 1(b).

TABLE I
MODIS BANDS USED IN THIS PAPER

Band	Bandwidth(μm)	Required NEΔT(K)
20	3.660 - 3.840	0.05
21	3.930 - 3.989	2.00
22	3.930 - 3.989	0.07
23	4.020 - 4.080	0.07
24	4.433 - 4.498	0.25
25	4.482 - 4.549	0.25
28	7.175 - 7.475	0.25
29	8.400 - 8.700	0.05
30	9.580 - 9.880	0.25
31	10.780 - 11.280	0.05
32	11.770 - 12.270	0.05
33	13.185 - 13.485	0.25
34	13.485 - 13.785	0.25
35	13.785 - 14.085	0.25

some type of transformed data. The image can be expressed as $\{\mathbf{x}(n), n = 1, 2, \ldots, N\}$, where $\mathbf{x}(n) = [x_1, x_2, \ldots, x_J]^T$. We may hypothesize that either the pixel $\mathbf{x}(n)$ consists of only noise $\mathbf{v}(n)$ and the background signal $\mathbf{b}(n)$ (the null hypothesis H_0), or that it also includes additional fire signal intensity $\mathbf{f}(n)$ (the H_1 hypothesis). Assuming that the signals are additive, we have

$$\mathbf{x}(n) = \mu \mathbf{f}(n) + \mathbf{b}(n) + \mathbf{v}(n). \tag{1}$$

Under hypothesis H_0, the fraction μ of the pixel on fire equals to 0; under hypothesis H_1, μ varies between 0 and 1.

If the background could be estimated, it can be subtracted from the original image to create a residual image. In the residual image, the two hypotheses are then simplified into the noise only hypothesis (H_0) and the fire plus noise hypothesis (H_1)

$$H_0 : \mathbf{x}(n) = \mathbf{v}(n) \tag{2}$$
$$H_1 : \mathbf{x}(n) = \mu \mathbf{f}(n) + \mathbf{v}(n) \tag{3}$$

Here $\mathbf{v}(n)$ is a noise only zero-mean process. A practical example of this background subtraction step and the properties of the residual image created are described in Section II.

Using the multivariate Gaussian model of the residual image in the local window, the conditional probability for a pixel to be under the null hypothesis H_0 is computed as

$$p(\mathbf{x} \mid H_0) = \frac{1}{(2\pi)^{J/2} |\Sigma|^{1/2}} \exp \left[-\frac{1}{2} \mathbf{x}^T \Sigma^{-1} \mathbf{x} \right]. \tag{4}$$

Here J is the number of features and $\mathbf{x}$ is the pixel under consideration. Σ is the unknown covariance matrix of the noise, which can be estimated using the residual image in the local window and (5).

$$\Sigma = E[\mathbf{x}\mathbf{x}^T] = \frac{1}{N} \Sigma_{n=1}^{N} \mathbf{x}(n)\mathbf{x}(n)^T \tag{5}$$

Σ is the maximum likelihood estimation of the unknown covariance matrix of noise under the null hypothesis H_0 (i.e., no fire). In Section V-B, we will describe a method to remove background fire pixels in order to estimate the noise covariance matrix for a local region.

We define the discriminant function $m(\mathbf{x})$ as the likelihood of H_0 with respect to $\mathbf{x}$, which is inversely proportional to the

squared Mahalanobis distance $\mathbf{x}^T \Sigma^{-1} \mathbf{x}$ [24]. So the discriminant function $m(\mathbf{x})$ can be expressed as follows:

$$m(\mathbf{x}) = \begin{cases} \mathbf{x}^T \Sigma^{-1} \mathbf{x} \geq D, & \text{fire pixel} \\ \mathbf{x}^T \Sigma^{-1} \mathbf{x} < D, & \text{background pixel}. \end{cases} \tag{6}$$

where D is the parameter of our algorithm that we need to tune manually using a small set of training images. Note that $m(\mathbf{x}) = \mathbf{x}^T \Sigma^{-1} \mathbf{x}$ is the squared Mahalanobis distance of the pixel being tested from the mean of background.

IV. IMAGE DATA

We test our algorithm using both satellite images and airborne images. Images from the MODIS sensors are used to characterize the performance of our algorithm on satellite data. MODIS has midwave infrared spectral bands specifically designed for fire detection and a number of other midwave and longwave infrared spectral bands suitable for detection of targets with a temperature range consistent with wildland fires. Fire detection with MODIS is enhanced with the aid of two midwave bands (band 21 and band 22) that measure radiance at the same spectral interval center at approximately 3.95 μm. However, band 22 has a saturation brightness temperature of about 330 K, while band 21 has lower gain and saturates at a brightness temperature of about 500 K. When band 22 is saturated, the value in band 21 is used instead, essentially extending the dynamic range of the system.

The airborne data is from the Wildfire Airborne Sensor Project at the Rochester Institute of Technology. WASP has been developed to provide a conceptual system for updating the existing U.S. Department of Agriculture Forest Service two band (MWIR and LWIR) Phoenix system for imaging fires [16]. Briefly, the WASP system contains 14-bit digital shortwave infrared (SWIR), MWIR, and LWIR cameras with bands centered at 1.3, 3.25, and 8.6 μm, respectively [15].

A. Feature Selection for the Modis Data

The MODIS bands used in our algorithm are listed in Table I. Many of the infrared bands listed in Table I are intended for atmospheric observations and are nominally within atmospheric absorption bands. However, we opt for including all of these 14 bands as features, because our algorithm can automatically exploit information in multiple bands and these bands typically

TABLE II
TEST IMAGES

Image Number	Date	Latitude Ranges	Longitude Ranges	Location
1	2003-07-18	60.0 ∼ 68.0	-180.0 ∼ -177.0	Siberia
2	2003-07-14	34.0 ∼ 40.0	-111.5 ∼ -108.0	Utah, USA
3	2003-08-10	40.0 ∼ 50.0	-117.0 ∼ -111.0	Idaho and Montana, USA
4	2002-10-19	-20.0 ∼ -10.0	138.0 ∼ 148.0	Northern Australia
5	2003-07-21	42.0 ∼ 50.0	-125.0 ∼ -118.0	Washington and Oregon, USA
6	2002-06-15	60.0 ∼ 66.5	-163.0 ∼ -149.5	Central Alaska, USA
7	2003-08-25	55.4 ∼ 63.5	77.7 ∼ 90.0	Southern Russia
8	2002-10-06	-10.4 ∼ -5.3	137.0 ∼ 145.1	New Guinea
9	2003-09-10	63.0 ∼ 69.0.0	147.0 ∼ 167.0	Siberia
10	2002-06-25	30.0 ∼ 36.5	-112.0 ∼ -106.0	Utah and Arizona, USA
11	2002-06-18	36.0 ∼ 40.0	-109.5 ∼ -103.5	Colorado and New Mexico, USA
12	2002-09-04	45.3 ∼ 51.8	90.8 ∼ 105.0	Northern Mongolia
13	2003-10-28	30 ∼38.0	-122.0 ∼ -113.5	Southern California, USA
14	2002-11-07	-37.0 ∼ -29.0	145.5 ∼ 154.0	New South Wales, Australia
15	2002-09-15	56.1 ∼ 63.0	122.1 ∼ 141.5	Yakutsk, Russia

contain some signal from fire. We do not include band 27 and band 36 in the feature set, because they are directly centered on strong atmospheric absorption bands. We find no images that exhibit fire signal in these two bands. All of the other bands are either in atmospheric windows or on the edges of atmospheric absorption bands and exhibit fire signal in enough cases to be useful, depending mostly on ground altitude and relative humidity. For example, for the Southern California images of October 28, 2003 shown in Fig. 6, the fires or parts of the fires (the hottest areas) are readily discernable in all MODIS thermal bands except 27 and 36. Incorporating these two bands only introduces noise without providing any useful information. Our experiments with and without these two bands show similar performance, which again indicates that these two bands contain little useful information. We therefore decide not to include them in the feature set.

Our algorithm uses one additional feature, the normalized thermal index, which is a transformation of the radiance data. The NTI was mainly developed for the detection of volcano eruptions, but its use for fire detection has also been discussed [12]. The NTI explores the nature of the relationship between temperature and radiance in the MWIR and LWIR based on differences in the slope of the Planck blackbody curve between these spectral regions. More specifically, the NTI is defined as

$$\text{NTI} = (L_{\text{MW}} - L_{\text{LW}})/(L_{\text{MW}} + L_{\text{LW}}) \qquad (7)$$

where, L_{MW} is the radiance received in the MWIR band and L_{LW} is the radiance received in the LWIR band. We use band 21 or band 22 for MWIR and band 32 for LWIR.

The MODIS images used in our experiments are listed in Table II. The first four images are from widely spaced geographic areas. We use them to train our algorithm in order to choose a proper value for the threshold D in (6). The rest of the images are used to test our algorithm for an extended variety of environments. The results shown below in image format are projected in order to provide geolocation information, but the algorithm is actually applied to the MODIS swath data.

B. Feature Selection for the WASP Data

For the WASP data, the feature set is the three infrared bands centered at 1.3, 3.25, and 8.6 μm, respectively, and NTI. The first WASP image was collected on April 17, 2004 over a prescribed burn in an oak-hickory forest in southern Ohio. The second one was taken on October 24, 2003, during an outdoor fire experiment in Rochester, NY.

V. IMPLEMENTATION DETAILS OF OUR ALGORITHM

In this section, we describe the detailed steps of our algorithm, including: 1) a prescreen step to eliminate obviously nonfire pixels; 2) a background characterization step to identify background fire pixels that should be excluded from the calculation of local statistical properties; and 3) a fire pixel identification step utilizing squared Mahalanobis distance to identify real fire pixels.

A. Identify Potential Fire Pixels

The first step of our algorithm is similar to other contextual fire detection algorithms [11]. To reduce the execution time of the algorithm, a nonfire prescreen method is performed to identify obviously nonfire pixels. To prevent real fire pixels from being classified as nonfire pixels, the thresholds in this step are very liberal. A daytime pixel is identified as a potential fire pixel only if the following two conditions are met. First, it must be satisfied that

$$(R_4 > \max\{\bar{R}_4, \bar{R}_{4_i}\}) \quad \text{and} \quad (R_{12} > \max\{\bar{R}_{12}, \bar{R}_{12_i}\})$$

where R_4 and R_{12} are the radiance in the 4-μm band and 12-μm band, respectively; $\bar{R}_4$ and $\bar{R}_{12}$ are the average radiance of the image in the two bands; and $\bar{R}_{4_i}$ and $\bar{R}_{12_i}$ are the means of the radiance in the two bands in the local window. For most images, this simple criterion can screen out more than half of the pixels to be processed. Moreover, it helps eliminate cloud and sun glint pixels. Reflective surfaces may appear warm in the 4-μm channel due to reflection but usually cool in the 12-μm channel, and clouds are usually cool in both the 4-μm channel and the 12-μm channel. The second condition that a potential fire pixel must meet is $\Delta T > 8$ K, where ΔT is the brightness temperature difference between the 4-μm channel and the 12-μm channel. This step is based on the significant increase in radiance at 4 μm relative to radiance at 12 μm for the active fire signature. Only those pixels that satisfy these two conditions are processed by the rest of the steps of our algorithm. The second condition of this prescreen step, however, is skipped

for the WASP images, because the WASP data are raw digital counts rather than radiance. The brightness temperature therefore cannot be calculated.

B. Background Characterization

The next step of our algorithm is to estimate the statistical properties of the neighboring background pixels in the local window centered on the potential fire pixel. The valid background pixels are those pixels in the local neighborhood that have no fire signal.

Most existing contextual algorithms (e.g., the MODIS algorithm) define background fire pixels as those with $T_4 > 325$ K and $T_{11} - T_4 > 20$ K for daytime observations, or those with $T_4 > 310$ K and $T_{11} - T_4 > 10$ K for night time observations. Unfortunately, this method performs dramatically different under different environments, because the thresholds are fixed and cannot adapt automatically. To remedy this limitation, we instead propose using the adaptive normalized thermal index to identify background fire pixels. More specifically, pixels that satisfy the following condition are flagged as background fire pixels and not considered in the calculation of the background statistics

$$\text{NTI} = (L_{\text{MW}} - L_{\text{LW}})/(L_{\text{MW}} + L_{\text{LW}}) > I. \qquad (8)$$

For MODIS images, we set the threshold $I = -0.64$; for WASP images, we set $I = -0.8$. The value of the NTI threshold for WASP data is different from that for MODIS data because of the difference in wavelength centers and bandpasses of the two different sensors, and more importantly, because the WASP images use raw digital counts rather than radiance. We determine these thresholds by examining images with known fire targets and set the thresholds sufficiently high to avoid mistakenly excluding warm background pixels as background fires. Although our use of NTI also requires the setting of a threshold, it is more adaptive than existing methods because it utilizes a normalized index and the threshold in NTI can be set with a big safety margin such that the result is much less sensitive to different environments when compared with existing methods.

Once the background fire pixels are masked, the multifeature image is demeaned using a 9×9 high-pass filter. Note that the background fire pixels are excluded when computing the local mean. Then the valid neighboring background pixels in the local window centered on the current potential fire pixel of interest are used to calculate the covariance matrix of the residual image. The window starts as a 31×31 pixel square region centered on the potential fire pixel. The window size is increased, if necessary, until at least 25% of the pixels in the window are valid background pixels. This approach assumes that the correlation between the fire pixel and the surrounding pixels decreases as the distance from the fire pixel increases.

C. Real Fire Identification Test

After feature selection and background characterization, the third step of our algorithm applies the maximum likelihood test in (6) to identify fire pixels. The maximum likelihood test determines if the sample under examination is a background pixel by testing whether the sample falls within a defined range of background variability, i.e., the squared Mahalanobis distance is less than a defined threshold.

The fire identification test requires the manual setting of a single threshold D. Before examining a set of images to determine the proper value for the threshold, we first need to define meaningful metrics that judge the quality of fire detection algorithms.

1) Evaluation Metric: Since we are working on a classification problem, we consider it appropriate to use *user accuracy* and *producer accuracy* to judge the performance of fire detection algorithms. Intuitively, user accuracy measures if the reported fires are real fires, while producer accuracy measures how many real fires are reported. These metrics are formally defined as

$$\text{user accuracy} = \frac{\text{\#correctly reported fires}}{\text{\#total reported fires}} \qquad (9)$$

$$\text{producer accuracy} = \frac{\text{\#correctly reported fires}}{\text{\#total fires that exist}}. \qquad (10)$$

It is important to note that these two metrics have to be used together, because it is always possible to improve one by sacrificing the other.

We decide not to use the popular *false alarm rate* metric, because the probability of fire in real images is very low. For example, suppose there are 10 fire pixels in a 450×400 pixel image. If all these fire pixels along with 10 false alarm pixels are detected, then the producer accuracy is $10/10 = 100\%$ while the false alarm rate will be $10/(450 * 400 - 10) = 0.005\%$. However, the user accuracy will be only $10/(10 + 10) = 50\%$. This example shows that, in the context of fire detection, the magnitude of user accuracy is more revealing than that of the false alarm rate.

Another subtle issue related to choosing a good metric for fire detection is the way to handle a single large fire that spans over multiple pixels. Most fire detection products report fires pixel by pixel. Only comparing pixel by pixel results, however, may lead to counterintuitive conclusions due to multiple detection from a large fire. For example, consider the case where a single large fire occupying a group of 20 adjacent pixels along with five discrete nonfire pixels are reported as fires by an algorithm and there is also a known single pixel fire that was not reported. The pixel method will have a user accuracy of 80% and a producer accuracy of 95%. That is overly optimistic. At the other extreme, one can treat a fire that spans over multiple adjacent pixels as a single object when computing user accuracy and producer accuracy, regardless of the size of the fire. With this method, the user accuracy and the producer accuracy for the algorithm would be 16.7% and 50%, respectively. This method, however, is too pessimistic in that it ignores the fact that the algorithm successfully finds a *big* fire. Here we propose counting the fires by fire regions but not by pixels, by using a sublinear damping function f.

For the metric to operate on a per fire region basis rather than a per pixel basis, a means for determining connected fire pixels is required. Because images of large fires may have some isolated fire pixels, we first use dilation [25] to join isolated pixels into nearby groups of fire pixels. We then counted the n reported fire
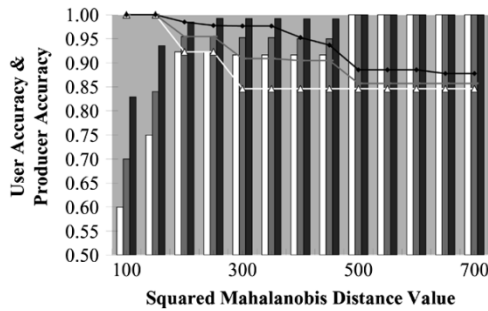
Fig. 3.   User accuracy (bars) and producer accuracy (lines) of our algorithm with respect to the squared Mahalanobis distance, with different damping functions: $f(n) = n$ (the black pair), $f(n) = \max(\sqrt{n}, 1)$ (the gray pair), and $f(n) = \max(\ln(n), 1)$ (the white pair).

pixels in the eight-connected region as $f(n)$ fires, where $f(n)$ is a damping function.

In Fig. 3, we compare three different functions: $f(n) = \max(\sqrt{n}, 1)$, $f(n) = \max(\ln(n), 1)$, and $f(n) = n$. The image of the large Rodeo–Chedeski fire in Arizona and two small fires in Mexico is used (the MODIS image 11). Since we do not know the exact fire pixels in the image, we use the fires reported by the MODIS fire product (obtained from EOS Data Gateway) as the ground truth. We run our algorithm for the image and report the user accuracy and producer accuracy metrics as a function of the squared Mahalanobis distance calculated in our algorithm [see (6)]. Fig. 3 shows that, as a function of the squared Mahalanobis distance, the shapes of the user accuracy and producer accuracy follow the same trend for different $f(n)$ functions, but the magnitude of the metrics are more reasonable under $f(n) = \max(\sqrt{n}, 1)$ and $f(n) = \max(\ln(n), 1)$ than that under the base case where $f(n) = n$. We opt for $f(n) = \max(\ln(n), 1)$ in the rest of our study as it produces metrics that are slightly more consistent with user intuition when we examine the fire images manually.

2) Threshold Determination: The detection threshold value was obtained empirically by testing our algorithm on the first four MODIS images listed in Table II, saving the rest of the images to assess the performance of the algorithm.

The ground truth of fires in the images are obtained in two ways. For images of locations in the U.S., the ground truth is obtained through the online GeoMAC internet-based mapping tool (http://geomac.usgs.gov/), which is designed for fire managers to access online maps of fire locations and perimeters in the U.S. For images of locations outside the United States, fires are identified by manual examination of the corresponding visible images at the resolution of 250×250 m. If smoke is observed in the visible image, the corresponding pixel is considered as a real fire pixel. We also considered the possibility of using high-resolution satellite images such as ASTER and Landsat images to get the ground truth. Unfortunately, due to the high-resolution of these images, we are unable to manually stitch together a large number of these images to get the ground truth of one MODIS image.

In Fig. 4, we vary the threshold D in (6) and observe its impact on user accuracy and producer accuracy. The test data are the first four MODIS images. We find that $D = 250$ is a value that strikes a good balance between these two metrics. In the
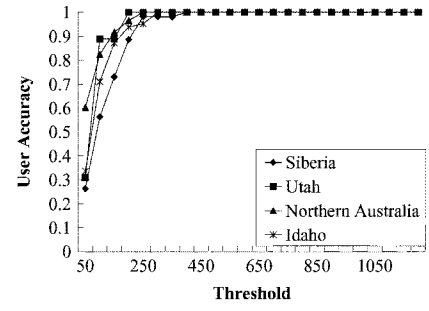
(a)
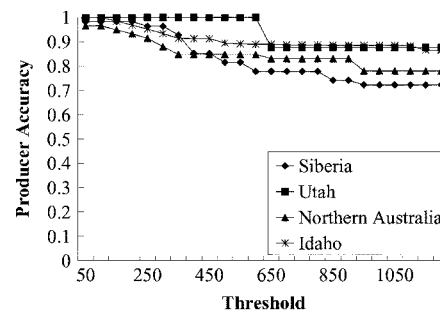
(b)

Fig. 4.   User accuracy and producer accuracy of our algorithm while varying the threshold D in (6). The test data are the first four MODIS images in Table II. $D = 250$ achieves a good balance between user accuracy and producer accuracy.

rest of this paper, we use this threshold for all the other MODIS images shown in Table II. We conduct a similar procedure for the WASP images and find that $D = 20$ is a good choice for the WASP images.

3) Fire Probability Image: Our algorithm can give a yes/no answer to the question whether a given pixel is a fire pixel. In addition, our novel use of squared Mahalanobis distance provides a direct means of deriving a fire probability image that is required for new approaches to fire propagation modeling. We denote by $P(\mathbf{x})$ the probability that a pixel $\mathbf{x}$ contain fire. $P(\mathbf{x})$ is only calculated for the potential fire pixels which passed the prescreen step described in Section V-A. Those nonfire pixels removed by the prescreen step are assigned 0% probability to contain fire. The function $P(\mathbf{x})$ should be a monotonic sublinear function that grows as the squared Mahalanobis distance $m(\mathbf{x})$ of pixel $\mathbf{x}$ increases. We define the probability function $P(\mathbf{x})$ as

$$P(\mathbf{x}) = 0.5 + 0.5 * \tanh\left(\frac{m(\mathbf{x}) - D}{D}\right) \qquad (11)$$

where D is the threshold in (6). With this probability function $P(\mathbf{x})$, the probability that pixel $\mathbf{x}$ contains a fire is 50% if the squared Mahalanobis distance of the pixel equals the threshold D.

VI. Experimental Results

A. Overall Performance

We first compare the overall performance of our algorithm with that of the latest MODIS fire product (version 4) using the MODIS test images in Table II. The ground truth to compute

TABLE III
PERFORMANCE COMPARISON BETWEEN THE HYBRID ALGORITHM AND MODIS FIRE PRODUCT

| Image Number	Location Name	Hybrid Algorithm		MODIS Fire Product	
		User Accuracy	Producer Accuracy	User Accuracy	Producer Accuracy
1	Siberia	0.9815	0.9636	1.0	0.9818
2	Utah, USA	1.0	1.0	1.0	1.0
3	Idaho and Montana, USA	0.9516	0.9672	0.9091	0.9677
4	Northern Australia	1.0	0.9138	0.9608	0.9423
5	Washington and Oregon, USA	1.0	0.875	1.0	0.875
6	Central Alaska, USA	1.0	1.0	0.8889	1.0
7	Southern Russia	0.9324	0.9324	1.0	0.7945
8	New Guinea	0.8429	0.8939	0.725	0.9355
9	Siberia	0.9620	0.9870	1.0	0.8701
10	Utah and Arizona, USA	0.9230	0.9230	1.0	1.0
11	Colorado and New Mexico, USA	1.0	0.9412	1.0	0.9412
12	Northern Mongolia	0.9130	1.0	0.9130	0.9545
13	Southern California, USA	1.0	1.0	0.9259	1.0
14	New South Wales, Australia	1.0	0.9444	0.9444	0.9444
15	Yakutsk, Russia	0.9605	0.9125	1.0	0.6951
	Average	0.9645	0.9502	0.9511	0.9268
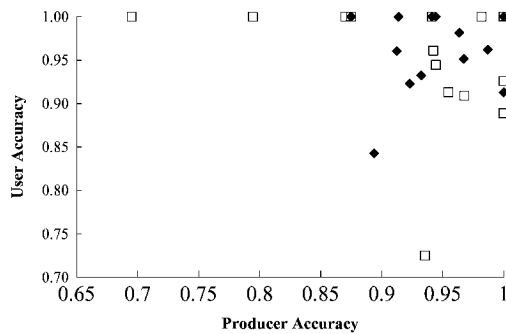
Fig. 5. User accuracy versus producer accuracy of the hybrid algorithm (diamonds) and MODIS fire product (hollow squares). The test images are listed in Table II.

TABLE IV
PERFORMANCE OF THE HYBRID ALGORITHM WHEN VARYING THE NUMBER OF MODIS CHANNELS USED AS FEATURES

Number of Channels	User Accuracy	Producer Accuracy
3	0.920	0.944
5	0.925	0.934
6	0.953	0.937
9	0.963	0.938
15	0.964	0.950
17	0.969	0.942

user accuracy and producer accuracy is from the GeoMAC data for images of locations in U.S., or from visual inspection for images of locations outside the U.S. The summary results are shown in Table III. We will present the detailed results for individual images in Section VI-B. Fig. 5 is a visual presentation of the results in Table III, i.e., the scatter plot of user accuracy versus producer accuracy for the individual.

In Fig. 5, both our algorithm and the MODIS fire product achieve a reasonable tradeoff between user accuracy and producer accuracy. On average, the user accuracy and the producer accuracy of our algorithm are about 1.5% and 2.6% higher than those of the MODIS fire product, respectively. In addition to this average performance improvement, more importantly, our algorithm achieves a better balance between user accuracy and producer accuracy for the images we tested. For the image of New Guinea, the user accuracy of the MODIS fire product is only 0.725; for the image of Yakutsk, the producer accuracy of the MODIS fire product is only 0.7125. By contrast, the user accuracy and the producer accuracy of our algorithm on all images are consistently higher than 80%. This performance shows that our simple algorithm with a single parameter is self-adaptive, i.e., its performance is not very sensitive to the particular image under test. This advantage of our algorithm over existing algorithms should be attributed to the sound statistical foundation behind our approach, and the new adaptive approach to identify background fire pixels.

The results presented above are from a configuration of our algorithm that uses the feature set described in Section IV-A, i.e., the 14 channels in Table I and the NTI from (7). We also perform experiments with other configurations of our algorithm using a varying number (3–17) of features. The results are shown in Table IV. For the three-channel configuration, we use bands 20, 21 or 22, and 32; for the five-channel configuration, we add bands 23 and 31; for the six-channel configuration, we further add NTI into the feature set; for the 9 channel configuration, we further add bands 24, 25, and 33; for the 15-channel configuration, we use the 14 channels in Table I and NTI ; and for the 17-channel configuration, we use all the 16 thermal bands of MODIS data and NTI.

The producer accuracy in Table IV is insensitive to the number of channels in use. The user accuracy improves slowly as the number of channels increases but even three bands already achieves good user accuracy and producer accuracy. Note that if we use only three bands in our algorithm, the MODIS fire product's user accuracy is slightly higher than that of our algorithm, although in all cases our algorithm has slightly higher producer accuracies. This is not surprising for us, since the MODIS fire product encodes the knowledge that researchers have accumulated over years to get rid of false alarms. On the contrary, our method does not have any additional rules to eliminate the false alarms. In Table IV, it is also shown that when we use more than six bands, our algorithm achieves both higher user accuracies and higher producer accuracies than the MODIS fire product for the images we tested. In many cases, the detected fires of the two algorithms appear to compliment each other. We believe that significant performance improvement over the state of the art can be obtained by combining the strength of our algorithm and the knowledge accumulated in
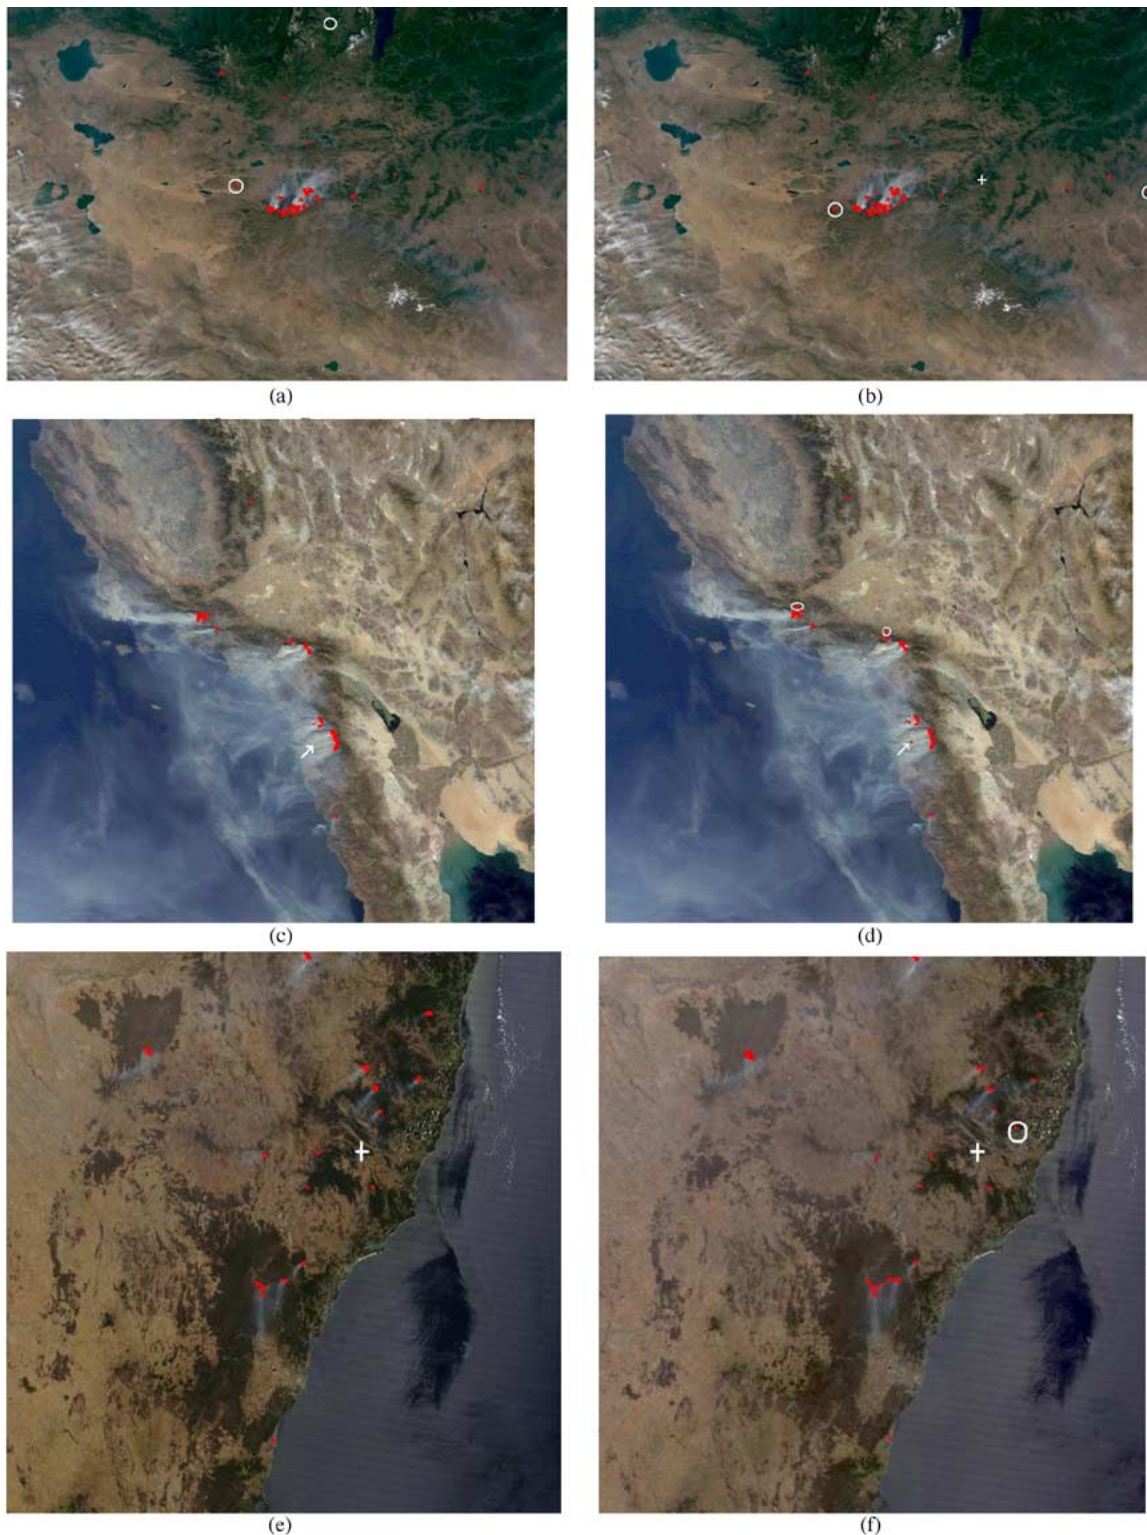
Fig. 6.   Results of the hybrid algorithm and MODIS fire products. The red dots represent the fires detected (which are enlarged), while white crosses indicate missed fires and white circles denote false alarms. (a) Result of the hybrid algorithm on image 12. (b) MODIS fire product of image 12. (c) Result of the hybrid algorithm on image 13. (d) MODIS fire product of image 13. (e) Result of the hybrid algorithm on image 14. (f) MODIS fire product of image 14.

the MODIS algorithm, i.e., building a full fire detection system that uses our algorithm as the core and then applies the rules in the MODIS algorithm to eliminate false alarms.

These results demonstrate two good properties of our algorithm. First, given a large number of bands, it can automatically utilize the important bands and the overall performance is

mostly influenced by those bands. This property greatly simplifies the tuning of the algorithm as we only need to set one global detection threshold D as opposed to one threshold for each band. Second, as bands that have less information are added into the feature set, our algorithm is able to automatically exploit the additional information in those bands to eliminate false alarms,
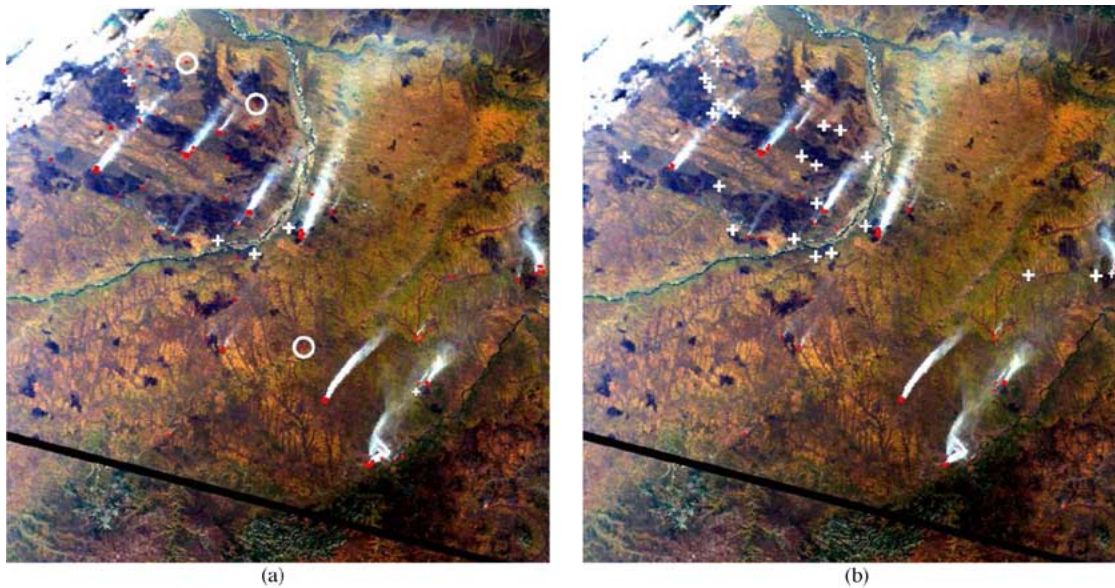
Fig. 7.   Results of the hybrid algorithm and MODIS fire products. The red dots represent the fires detected (which are enlarged), while white crosses indicate missed fires and white circles denote false alarms. (a) Result of the hybrid algorithm on image 15. (b) MODIS fire product of image 15.

as is indicated by the improvement in user accuracy. Both good properties of our algorithm should be attributed to its sound statistical model that treats fire detection as a statistical anomaly detection problem.

In order to test the flexibility of our algorithm, we also applied it to two airborne images taken by the WASP sensor. The bands used in the WASP images are different from those used in the MODIS images. For both WASP images, our algorithm successfully finds all fires without any false alarms. We will present visualized results for these images later. The successes of our algorithm on images taken from different types of sensors demonstrate that our algorithm is truly adaptive.

B. Performance for Individual Images

Next, we proceed to show the visualized results for individual images (the MODIS test images 12–15, and the two WASP images) in Figs. 6–8. In all figures, the red dots, white crosses, and white circles represent detected fires, missed fires, and false alarms, respectively.

Fig. 6(a) and (b) shows result of the Mongolia image (image 12 in Table II produced by our algorithm and the MODIS fire product, respectively. A careful visual examination of the corresponding images in the 250-m visible image, midinfrared, and long wave infrared bands reveals that our algorithm reports two separate false alarms (indicated with white circles). The MODIS fire product also reports two false alarms, but at different locations. In addition, it also misses a small fire, which is accompanied by an obvious smoke plume in the 250-m visible image.

Fig. 6(c) and (d) shows results for an image of Southern California taken on October 28, 2003 (image 13 in Table II) when large wild fires engulfing vast area of Southern California. There are five groups of fires in the image. Below we describe them in the order from top to bottom. The small single-pixel on the top is reported by both our algorithm and the MODIS fire product. Although we are unable to find information for this fire in the GeoMAC (http://geomac.usgs.gov/) reported fire database, the

smoke plume for this fire can be clearly seen in the 250-m visible image. Therefore, we consider this fire to be a true fire for both algorithms rather than a false alarm. The second group of fires down from the top consist of the Piru Fire, the Verdale Fire, and the Simi Incident Fire in the GeoMAC database. The third group of fires down from the top is a combination of the Grand Prix Fire, the Padus Fire, and the Old Fire. The fourth group of fires down from the top (fires close to the white arrow) consist of the Paradise Fire, and the Cedar Fire. Last, the small fire at the bottom is reported by both algorithms. Because the location of this fire is in Mexico, the GeoMAC database has no information for this fire. However, we consider it as a true fire because there is clearly visible smoke plume in the image. In Fig. 6(d), the MODIS fire product reports two false alarms (indicated by the white circles), neither of which has sufficiently high brightness temperature in the 4-μm channel to suggest the present of a true fire. In addition, for these reported fire pixels, no smoke plumes can be seen in the visible image. In Fig. 6(d), the MODIS fire product also reports a small fire (marked by the white arrow) that is absent in the result of our algorithm. The location of this fire is on the edge of the burn scar of the Cedar Fire. It is hard to judge whether this pixel contains a real fire since the area is completely covered by heavy smoke possibly generated by the adjacent big fires. Therefore, it may be a small extinguishing fire or a recently burned patch that is still hot. Due to this difficulty, we ignore this one-pixel fire in the comparison between our algorithm and the MODIS fire product, i.e., not counting it when calculating user accuracy and producer accuracy.

Fig. 6(e) and (f) shows an image covering a portion of New South Wales, Australia (image 14 in Table II). In November 2002, warm and dry weather led to at least 55 separate wild fires burning across New South Wales. In the image taken on November 7, we can see 16 separate fires indicated by heavy smokes in the image. Our algorithm misses a small fire but reports no false alarms. All fires reported by our algorithm are accompanied by discernible smoke in the visible image. The
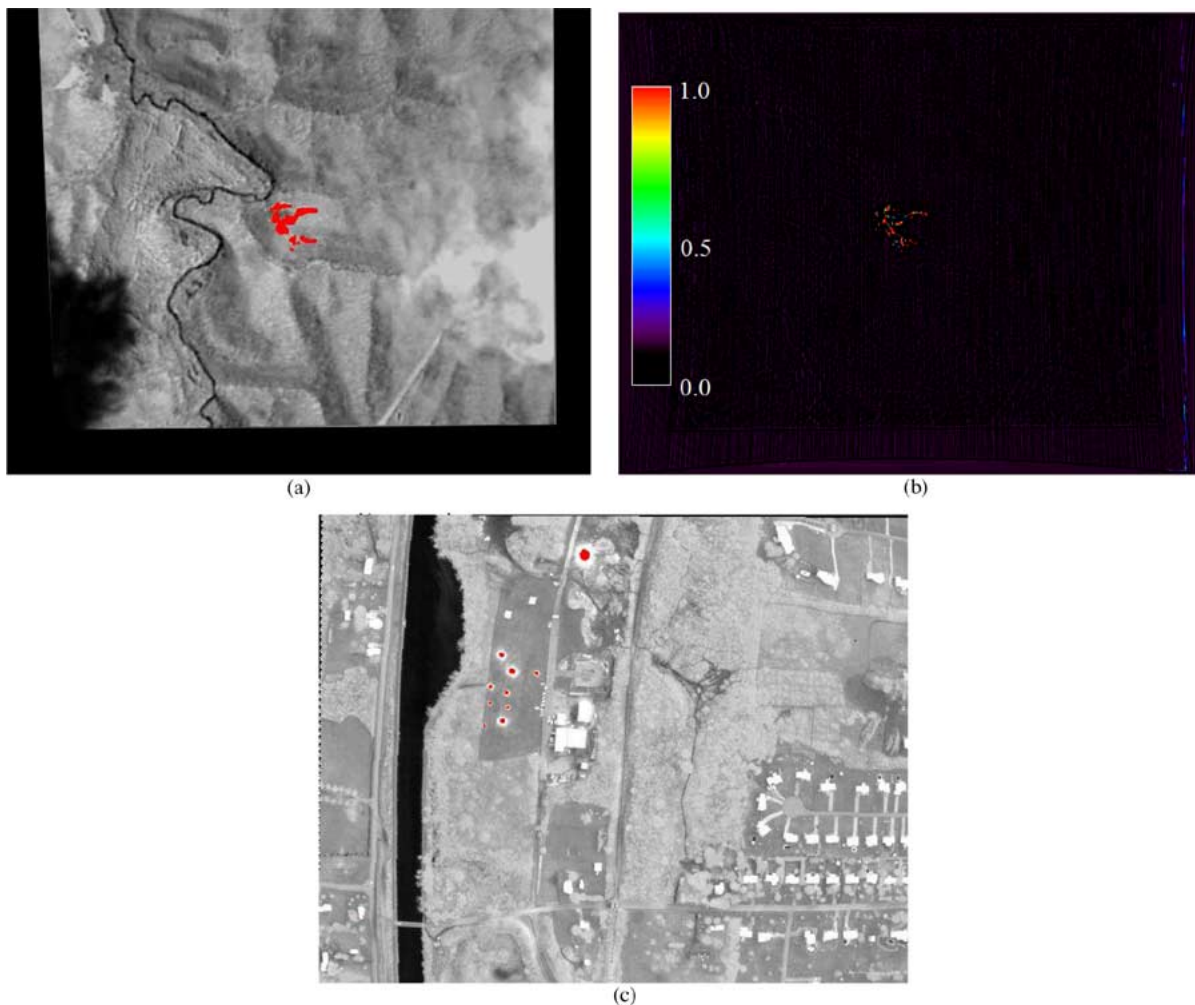
Fig. 8.   Results of the hybrid algorithm on test images. The red dots represent fires detected by the hybrid algorithm (the red dots are enlarged to be seen). (a) Result of the Hybrid algorithm on the WASP image of a prescribed fire in Ohio (b) The fire probability image produced by the Hybrid algorithm. (c) Result of the Hybrid algorithm on the WASP image of Spencerport area with prescribed fires.

MODIS fire product misses the same fire that our algorithm misses. In addition, it also reports a false alarm, for which no discernable smoke plume can be seen in the high-resolution visible image.

Fig. 7(a) and (b) shows an image of Yakutsk, in Russia (image 15 in Table II). Several big fires with heavy smoke plumes as well as many small and cool fires were burning when the image was taken on September 15, 2002. The region in this image is very cold. The average brightness temperature of this image (including fire pixels) is only 302 K. Our algorithm misses six small, cool fires, and mistakenly reports three false alarms. By contrast, the MODIS fire product introduces no false alarms but misses 25 small fires that can be clearly identified by the visible smoke plumes in the image.

Our algorithm automatically identifies background pixels using a statistical approach. Because of this flexibility, it can work with images produced by different thermal sensors with different band sets. We demonstrate this capability of our algorithm on two WASP images that we collected over prescribed fires.

Fig. 8(a) is a WASP image captured on April 17, 2004 over a prescribed fire in an oak-hickory forest in southern Ohio. The

WASP sensor has a ground spot diameter of 3 m and was flown at a 3048-m altitude. The red dots in Fig. 8(a) are detected fires, which precisely match the fire location of this experiment. Fig. 8(b) presents the result in Fig. 8(a) in a different manner. Our algorithm's confidence level of fire presenting in a given pixel is represented by different colors (see the left of the image), with black and red representing probability zero and one, respectively. The probability is calculated from (11). In this figure, our algorithm clearly separates background pixels from fire pixels, with a smooth transition probability from 0 to 1 at the edge of the fire. In addition to the yes/no answer that our algorithm provides, this probability figure provides additional information on the confidence level of our algorithm. Since no fire detection algorithm can be perfect for all situations, this confidence information can be very helpful in guiding users to make action decision.

Fig. 8(c) is the result of our algorithm on another WASP image, which was taken on October 24, 2003 during a prescribed outdoor fire experiment. The image is obtained at the altitude of 1219 m, and the corresponding pixel size is 1.2×1.2 m. Ten fire targets of different sizes and created with different kinds of fuel were set up. Our algorithm successfully detected all the

fires without any false alarm. These results again indicate that our algorithm does not depend on specific band sets, and can be easily transplanted to different thermal sensors. We do not include the fire probability image for this WASP image. Because the fire targets in this image are very small (most of them are subpixel targets), these one pixel fire targets in the probability image are invisible without enlargement.

VII. CONCLUSION

In this paper, we proposed a hybrid contextual fire detection algorithm that leverages the well-developed contextual approach to address variation in image background. Our algorithm uses a multivariate distance metric (the squared Mahalanobis distance) operating in the feature space to discriminate fire pixels. This feature-based approach allows us to incorporate a larger number of spectral bands and transformations such as the normalized thermal index. Compared with the latest MODIS fire product (version 4), our algorithm improves both user accuracy and producer accuracy by 1.5% and 2.6% on average, respectively, and up to 28% for some images. In many cases, the fires detected by the two algorithms appear to complement each other, suggesting that combining the algorithms may be useful. We leave this as a subject of future research.

Our algorithm has several major advantages over existing contextual algorithms. First, our algorithm is adaptive and flexible. The results shown in Table IV indicate that our algorithm can still work well when some bands are not available due to sensor problems or other causes by using the predetermined threshold for different feature number. Further, it can be easily applied to different sensors that use different band sets, as is demonstrated by the excellent performance with images taken by the MODIS sensors and our own WASP sensor. Second, our algorithm is relatively simple. It includes no manually constructed rules and relies on only two insensitive parameters (one for identifying background fire pixels and the other one for detecting real fires) that can be tuned easily through a very small training set (four images). Third, our feature-based algorithm does not need the calibrated radiance as the input, extending the utility of the method to operational sensors that may not have the laboratory capabilities for producing radiance calibrations. Finally, our novel use of the squared Mahalanobis distance metric provides a means for producing a pixel by pixel fire probability image, which can be an important input to statistical methods for fire propagation estimation.

REFERENCES

[1] P. J. Crutzen, L. E. Heidt, J. P. Krasnec, W. H. Pollock, and W. Seiler, "Biomass burning as a source of atmospheric gases, CO, H$_2$, N$_2$O, NO, CH$_3$Cl, and COS," *Nature*, vol. 282, pp. 253–256, 1979.

[2] P. Crutzen and M. O. Andreae, "Biomass burning in the tropics: Impact on atmospheric chemistry and biogeochemical cycle," *Science*, vol. 250, pp. 1669–1678, 1990.

[3] Y. J. Kaufman and R. S. Fraser, "The effect of smoke particles on clouds and climate forcing," *Science*, vol. 277, pp. 1636–1639, 1997.

[4] Z. Li, "Influence of absorbing aerosols on the inference of solar surface radiation budget and cloud absorption," *J. Clim.*, vol. 11, pp. 5–17, 1998.

[5] R. A. Houghton and G. M. Woodwell, "Global climatic change," *Sci. Amer.*, vol. 260, pp. 36–44, 1989.

[6] H. A. Mooney, P. M. Vitousek, and P. A. Matson, "Exchange of materials between terrestrial ecosystems and the atmosphere," *Science*, vol. 238, pp. 926–932, 1987.

[7] T. L. Clark, J. L. Coen, and D. Latham, "Description of a coupled atmosphere-fire model," *Int. J. Wildland Fire*, vol. 13, pp. 2783–2799, 1992.

[8] J. Mandel, M. Chen, J. L. Coen, C. C. Douglas, L. Franca, C. Johns, R. Kremens, A. Puhalskii, A. Vodacek, and W. Zhao, "Dynamic data driven wildfire modeling," in *Dynamic Data Driven Applications Systems*, F. Darema, Ed. Amsterdam, The Netherlands: Kluwer, 2005, to be published.

[9] J. Dozier, "A method for satellite identification of surface temperature fields of subpixel resolution," *Remote Sens. Environ.*, vol. 11, pp. 221–229, 1981.

[10] Y. J. Kaufman, C. O. Justice, L. P. Flynn, J. D. Kendall, E. M. Prins, L. Giglio, D. E. Ward, W. P. Menzel, and A. W. Setzer, "Potential global fire monitoring from EOS-MODIS," *J. Geophys. Res.*, vol. 103, pp. 32 215–32 238, 1998.

[11] L. Giglio, J. Descloitres, C. O. Justice, and Y. J. Kaufman, "An enhanced contextual fire detection algorithm for MODIS," *Remote Sens. Environ.*, vol. 87, pp. 273–282, 2003.

[12] R. Wright, L. Flynn, H. Garbeil, A. Harris, and E. Pilger, "Automated volcanic eruption detection using MODIS," *Remote Sens. Environ.*, vol. 82, pp. 135–155, 2002.

[13] L. R. Radke, T. L. Clark, J. L. Coen, C. Walther, R. N. Lockwood, P. J. Riggan, J. Brass, and R. Higgins, "The Wildfire Experiment (WiFE): Observations with airborne remote sensors," *Can. J. Remote Sens.ing*, vol. 26, pp. 406–417, 2000.

[14] V. G. Ambrosia, S. S. Wegener, C. V. Sullivan, S. W. Buechel, S. E. Dunagan, J. A. Brass, and J. Stoneburner, "Demonstrating UAV-acquired real-time thermal data over fires," *J. Amer. Soc. Photogramm. Remote Sens.*, vol. 69, pp. 391–402, 2003.

[15] R. L. Kremens, J. Faulring, D. McKeown, M. Richardson, J. Cock-burn, G. Semeraro, and H. Rhody, "Wildfire airborne sensor project (WASP)," *Photogramm. Eng. Remote Sens.*, 2004, submitted for publication.

[16] P. H. Greenfield, W. Smith, and D. C. Chamberlain, "Phoenix-the new forest service airborne infrared fire detection and mapping system," presented at the *2nd Int. Wildland Fire Ecology and Fire Management Congr. and the 5th Symp. Fire and Forest Meteorology*, Orlando, Florida, 2003.

[17] E. M. Prins and W. P. Menzel, "Geostationary satellite detection of biomass burning in South America," *Intl. J. Remote Sens.*, vol. 13, pp. 49–63, 2004.

[18] P. J. Kennedy, A. S. Belward, and J. M. Gregoire, "An improved approach to fire monitoring in West Africa using AVHRR data," *Int. J. Remote Sens.*, vol. 15, pp. 2235–2255, 1994.

[19] D. Manolakies, C. Siracusa, D. Marden, and G. Shaw, "Hyperspectral adaptive matched filter detectors: Practical performance comparison," in *Proc. SPIE*, vol. 4381, 2001, pp. 18–33.

[20] D. Kuan, A. Sawchuk, T. Strand, and P. Chavel, "Adaptive noise smoothing filter for images with signal-dependent noise," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. PAMI-7, no. 2, pp. 165–177, 1985.

[21] B. R. Hunt and T. M. Cannon, "Nonstationary assumptions of Gaussian models of images," *IEEE Trans. Syst., Man, Cybern.*, vol. SMC-6, pp. 876–881, 1976.

[22] I. S. Reed and X. Yu, "Adaptive multiple-band CFAR detection of an optical pattern with unknown spectral distribution," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 38, no. 10, pp. 1760–1770, Oct. 1990.

[23] A. Margalit, I. Reed, and R. M. Gagliardi, "Adaptive optical target detection using correlated images," *IEEE Trans. Aerosp. Electron. Syst.*, vol. AES-21, 1985.

[24] R. J. Muirhead, *Aspects of Multivariate Statistical Theory*, 1st ed. New York: Wiley, 1982.

[25] C. G. Rafael and E. Richard, *Digital Image Processing*, 3rd ed. Reading, MA: Addison-Wesley, 1992.

Ying Li received the B.S. and M.E. degrees from the University of Science and Technology of China, Hefei, in 1997, and 2000, respectively. She is currently pursuing the Ph.D. degree at the Rochester Institute of Technology, Rochester, NY.

She is now working on wildland fire detection and mapping using satellite/airborne images. Her research interests lie in the area of digital image processing, pattern classification, and image segmentation.

Anthony Vodacek received the B.S. degree in chemistry from the University of Wisconsin-Madison, and the M.S. and Ph.D. degrees in environmental engineering from Cornell University, Ithaca, NY, in 1981, 1985, and 1990, respectively.
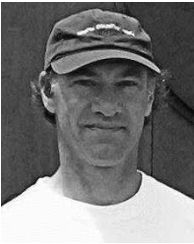
He has held Postdoctoral positions at the Joint Research Centre in Ispra, Italy and at Goddard Space Flight Center, Wallops Flight facility. He then held research positions at the University of Maryland in the Department of Chemistry and Biochemistry, before joining the faculty in the Center for Imaging Science, Rochester Institute of Technology, Rochester, NY, where he is now Associate Professor. His research interests include spectroradiometric measurements and spectral processing methods for images of aquatic and terrestrial environments.

Robert L. Kremens received the B.S. degree in physics from Cooper Union, New York, the M.S. and Ph.D. degrees in physics from New York University, New York, and the M.S. degree in environmental studies from the University of Rochester, Rochester, NY, in 1975, 1977, 1981, and 2000, respectively.

He has been with the Rochester Institute of Technology (RIT), Rochester, at the Center for Imaging Science since 2000. At RIT, he has specialized in physical measurements of wildland fires for remote sensing and fire behavior purposes and in constructing new airborne and ground-based systems to monitor the environment and wildland fire. Before RIT, he worked at the University of Rochester on the world's largest pulsed laser, building and analyzing nuclear and laser diagnostics on an inertial fusion experiment, at the U.S. Army Ballistics Research Laboratory, Aberdeen, MD, and at several industrial positions where he designed and constructed high-speed data acquisition and imaging systems.

Ambrose Ononye received the M.S. degree in physics and the Ph.D. degree in electrical engineering.

He is a Postdoctoral Research Scientist at the Center of Imaging Science, Rochester Institute of Technology, Rochester, NY. His research interests are in the areas of digital image processing, remote sensing, and computer vision.

Chunqiang Tang (S'03–M'04) received the B.E. degree from the University of Science and Technology of China, Hefei, in 1996, the M.S. degree from the Institute of Computing Technology, Chinese Academy of Sciences, Beijing, in 1999, and the Ph.D. degree from the University of Rochester, Rochester, NY, in 1996, 1999, and 2004, respectively.

He joined the IBM T. J. Watson Research Center, Yorktown Heights, NY, as a Research Staff Member in August 2004. His research interests lie in the areas of information retrieval, distributed systems, computer networks, and operating systems.

Dr. Tang is a member of ACM.