

Fire detection in infrared video using wavelet analysis

Behcet Uğur Töreysin
Ramazan Gökberk Cinbiş
Yiğithan Dedeoğlu
Ahmet Enis Çetin
Bilkent University
TR-06800 Bilkent
Ankara, Turkey

Abstract. A novel method to detect flames in infrared (IR) video is proposed. Image regions containing flames appear as bright regions in IR video. In addition to ordinary motion and brightness clues, the flame flicker process is also detected by using a hidden Markov model (HMM) describing the temporal behavior. IR image frames are also analyzed spatially. Boundaries of flames are represented in wavelet domain and the high frequency nature of the boundaries of fire regions is also used as a clue to model the flame flicker. All of the temporal and spatial clues extracted from the IR video are combined to reach a final decision. False alarms due to ordinary bright moving objects are greatly reduced because of the HMM-based flicker modeling and wavelet domain boundary modeling. © 2007 Society of Photo-Optical Instrumentation Engineers.
[DOI: 10.1117/1.2748752]

Subject terms: infrared video fire detection; wavelet transform; video object contour analysis; video event detection; segmentation; hidden Markov models.

Paper 060481R received Jun. 22, 2006; revised manuscript received Dec. 1, 2006; accepted for publication Dec. 19, 2006; published online Jun. 29, 2007.

1 Introduction

Conventional point smoke and fire detectors typically detect the presence of certain particles generated by smoke and fire by ionization or photometry. An important weakness of point detectors is that they cannot provide quick responses in open or large spaces. In this work, an automatic fire detection algorithm in infrared (IR) video is proposed. The strength of using video in fire detection is the ability to monitor large and open spaces such as auditoriums and atriums. Recently, fire and flame detection algorithms in regular video were developed.¹⁻⁵ They are based on the use of color and motion information in video. Current algorithms are not robust in outdoor applications, for example, they may produce false alarms to reddish leaves flickering in the wind and reflections of periodic warning lights. Other recent methods for video-based fire detection are Refs. 6-8. These methods are developed to detect the presence of smoke in the video.

It was reported that turbulent flames flicker with a frequency of around 10 Hz.^{2,9,10} Various flame flicker values were reported for different fuel types in Refs. 11 and 12. In practice, the flame flicker process is time varying and it is far from periodic. This stochastic behavior in flicker frequency is especially valid for uncontrolled fires. Therefore, a random model-based modeling of the flame flicker process produces more robust performance compared to frequency-domain-based methods that try to detect peaks around 10 Hz in the Fourier domain. In Ref. 4, fire and flame flicker is modeled by using hidden Markov models (HMMs) in visible video. The use of IR cameras instead of a regular camera provides further robustness to imaging-based fire detection systems, especially for fires with little radiance in the visible spectrum, e.g., alcohol and hydrogen

fires, which are common in tunnel collisions. Unfortunately, the algorithms developed for regular video cannot be used in IR video due to the lack of color information, and there is almost no spatial variation or very little texture information in fire regions in IR video as in most hot objects. Therefore, new image analysis techniques have to be developed to automatically detect fire in IR videos.

A bright-looking object in IR video exhibiting rapid time-varying contours is an important sign of presence of flames in the scene. This time-varying behavior is not only directly observable in the contours of a fire region, but is also observable as variations of color channel values of the pixels in regular video. On the other hand, we do not observe this behavior in IR videos. The entire fire region appears as a flat white region in IR cameras operating in white-hot mode.

In this work, boundaries of moving bright regions are estimated in each IR image frame. A 1-D curve representing the distance to the boundary from the center of mass of the region is extracted for each moving hot region. The wavelet transform of this 1-D curve is computed and the high frequency nature of the contour of the fire region is determined using the energy of the wavelet signal. This spatial domain clue is also combined with temporal clues to reach a final decision.

The organization of this work is as follows. In Sec. 2.1, the spatial image analysis and feature extraction method based on wavelet analysis is described. In Sec. 2.2, temporal video analysis and HMM-based modeling of the flicker process is presented. Simulation examples are presented in Sec. 3.

2 Fire and Flame Behavior in Infrared Video

Fire and flame detection methods in regular video use color information.^{1-5,13} On the other hand, most IR imaging sensors provide a measure of the heat distribution in the scene

and each pixel has a single value. Usually, hot (cold) objects in the scene are displayed as bright (dark) regions in white-hot mode in IR cameras. Therefore, fire and flame pixels appear as local maxima in an IR image. If a relatively bright region moves in the captured video, then it should be marked as a potential region of fire in the scene monitored by the IR camera. However, an algorithm based on only motion and brightness information will produce many false alarms because vehicles, animals, and people are warmer than the background, and they also appear as bright objects. In this work, in addition to motion and relative brightness information, object boundaries are analyzed both spatially (intraframe) and temporally (interframe).

Boundaries of uncontrolled fire regions in an image frame are obviously irregular. On the other hand, almost all regular objects and people have smooth and stationary boundaries. This information is modeled using wavelet domain analysis of moving object contours, which is described in the next section. One can reduce the false alarms that may be due to ordinary moving hot objects by carrying out temporal analysis around object boundaries to detect random changes in object contours. This analysis is described in Sec. 2.2.

2.1 Wavelet Domain Analysis of Object Contours


Moving objects in IR video are detected using the background estimation method developed by Collins, Lipton, and Kanade.¹⁴ This method assumes that the camera is stationary. Moving pixels are determined by subtracting the current image from the background image and thresholding. A recursive adaptive threshold estimation is described in Ref. 14 as well. Other methods can be also used for moving object estimation. After moving object detection, it is checked whether the object is hotter than the background, i.e., it is verified if some of the object pixels are higher in value than the background pixels.

Hot objects and regions in IR video can be determined in moving cameras as well by estimating local maxima in the image. Contours of these high temperature regions can be determined by region growing.


The next step of the proposed method is to determine the center of mass of the moving bright object. A 1-D signal $x(\theta)$ is obtained by computing the distance from the center of mass of the object to the object boundary for $0 \leq \theta < 2\pi$. In Fig. 1, two forward-looking infrared (FLIR) image frames are shown. The example feature functions for the walking man, indicated with an arrow, and the fire region in Fig. 1 are shown in Fig. 2 for 64 equally spaced angles $x[l] = x(l\theta_s)$, $\theta_s = 2\pi/64$. To determine the high-frequency content of a curve, we use a single scale wavelet transform shown in Fig. 3. The feature signal $x[l]$ is fed to a filter bank shown in Fig. 3 and the low-band signal

$$c[l] = \sum_m h[2l - m]x[m], \quad (1)$$

and the high-band subsignal



(a)



(b)

Fig. 1 Two relatively bright moving objects in FLIR video: (a) fire image and (b) a man (shown with an arrow). Moving objects are determined by the hybrid background subtraction algorithm of Ref. 14.

$$w[l] = \sum_m g[2l - m]x[m], \quad (2)$$

are obtained. Coefficients of the low-pass and high-pass filters are $h[l] = \{1/4, 1/2, 1/4\}$ and $g[l] = \{-1/4, 1/2, -1/4\}$, respectively.¹⁵⁻¹⁷

The absolute values of high-band (wavelet) $w[l]$ and low-band $c[l]$ coefficients of the fire region and the walking man are shown in Figs. 4 and 5, respectively. The high-frequency variations of the feature signal of the fire region are clearly distinct from that of the man. Since regular objects have relatively smooth boundaries compared to flames, the high-frequency wavelet coefficients of flame boundary feature signals have more energy than regular objects. Therefore, the ratio of the wavelet domain energy to the energy of the low-band signal is a good indicator of a fire region. This ratio is defined as

$$\rho = \frac{\sum_l |w[l]|}{\sum_l |c[l]|}. \quad (3)$$

The likelihood of the moving region to be a fire region is highly correlated with the parameter ρ . The higher the




Fig. 2 Equally spaced 64 contour points of (a) the walking man and (b) the fire regions shown in Fig. 1.

value of ρ , the higher the probability of the region belonging to flame regions.

A threshold ρ_T for ρ was experimentally estimated offline. During real-time analysis, regions for which $\rho > \rho_T$ are first determined. Such regions are possible fire regions. In order not to miss any fire region, a low threshold value for ρ_T is selected. Therefore, temporal flicker analysis should be carried out in these regions to reach a final decision. The flicker detection process is described in the next section.

2.2 Modeling Temporal Flame Behavior

It was reported in mechanical engineering literature that turbulent flames flicker with a frequency of 10 Hz.¹⁰ In Ref. 18, the shape of fire regions are represented in the Fourier domain. Since Fourier transform does not carry any time information, fast Fourier transforms (FFTs) have to be computed in windows of data and temporal window size, and the peak or energy around 10 Hz is very critical for flicker detection. If the window size is too long, then one may not get enough peaks in the FFT data. If it is too short,




Fig. 3 Single-stage wavelet filter bank. The high-pass and low-pass filter coefficients are $\{-1/4, 1/2, -1/4\}$ and $\{1/4, 1/2, 1/4\}$, respectively.

then one may completely miss flicker and therefore no peaks can be observed in the Fourier domain. Furthermore, one may not observe a peak at 10 Hz but a plateau around it, which may be hard to distinguish from the Fourier domain background.

Another problem is that one may not detect periodicity in fast growing uncontrolled fires, because the boundary of the fire region simply grows in video. Actually, the fire behavior is a wide-band random activity below 15 Hz, and a random-process-based modeling approach is naturally suited to characterize the rapid time-varying characteristic of flame boundaries. Broadbent and Huang et al. independently reported different flicker frequency distributions for various fuel types in Refs. 11 and 12. In general a pixel, especially at the edge of a flame, becomes part of the fire and disappears in the background several times in one second of a video at random. In fact, we analyzed the temporal characteristics of the red channel value of a pixel at the boundary of a flame region in color-video clips recorded at 10 and 25 fps. We also analyzed the temporal characteristics of the intensity value of a pixel at the boundary of a flame region in an IR video clip recorded at 10 fps. We obtained the flicker frequency distributions shown in Fig. 6 for 10 fps color video, 25 fps color video, and 10 fps IR video, respectively. We assumed that the flame flicker behavior is a wide-band random activity below 15 Hz for all practical purposes. This is the basic reason behind our stochastic model.

Flame flicker can be detected in low-rate image sequences obtained with a rate of less than 20 Hz as well, in spite of the aliasing phenomenon. To capture 10-Hz flicker, the video should capture at least 20 fps. However, in some surveillance systems, the video capture rate is below 20 Hz. If the video is available at a lower capture rate, aliasing




Fig. 4 The absolute values of (a) high-band (wavelet) and (b) low-band coefficients for the fire region.

occurs but flicker due to flames can still be observed in the video (Fig. 6). For example, 8-Hz sinusoid appears as 2-Hz sinusoid in a 10 fps video.⁵ An aliased version of flame flicker signal is also a wide-band signal in the discrete-time Fourier transform domain. This characteristic flicker behavior is very well suited to be modeled as a random Markov

model, which is extensively used in speech recognition systems, and recently these models have been used in computer vision applications.¹⁹

In this work, three-state Markov models are trained offline for both flame and nonflame pixels to represent the temporal behavior (Fig. 7). These models are trained using




Fig. 5 The absolute (a) high-band (wavelet) and (b) low-band coefficients for the walking man.




Fig. 6 Flicker frequency distributions for (a) 10-fps color video, (b) 25-fps color video, and (c) 10-fps IR video. These frequency distributions were obtained by analyzing the temporal variations in the red channel value of a pixel at the boundary of a flame region in color-video clips recorded at 10 and 25 fps, and intensity value of a pixel at the boundary of a flame region in an IR video clip recorded at 10 fps, respectively.

a feature signal, which is defined as follows: let $I_k(n)$ be the intensity value of the k 'th pixel at frame n . The wavelet coefficients of I_k are obtained by the same structure shown in Fig. 3, but filtering is implemented temporally.

Wavelet signals can easily reveal the random characteristic of a given signal, which is an intrinsic nature of flame pixels. That is why the use of wavelets instead of actual pixel values leads to more robust detection of flames in video. Since wavelet signals are high-pass filtered signals, slow variations in the original signal lead to zero-valued wavelet coefficients. Hence it is easier to set thresholds in the wavelet domain to distinguish slow varying signals from rapidly changing signals. Non-negative thresholds $T_1 < T_2$ are introduced in the wavelet domain to define the three states of the hidden Markov models for flame and nonflame moving bright objects.




Fig. 7 Three-state Markov models for (a) flame and (b) nonflame moving pixels.

The states of HMMs are defined as follows: at time n , if $|w(n)| < T_1$, the state is in $S1$; if $T_1 < |w(n)| < T_2$, the state is $S2$; else if $|w(n)| > T_2$, the state $S3$ is attained. For the pixels of regular hot objects like walking people, the engine of a moving car, etc., no rapid changes take place in the pixel values. Therefore, the temporal wavelet coefficients ideally should be zero, but due to thermal noise of the camera, the wavelet coefficients wiggle around zero. The lower threshold T_1 basically determines a given wavelet coefficient being close to zero. The state defined for the wavelet coefficients below T_1 is $S1$. The second threshold T_2 indicates that the wavelet coefficient is significantly higher than zero. The state defined for the wavelet coefficients above this second threshold T_2 is $S3$. The values between T_1 and T_2 define $S2$. The state $S2$ provides hysteresis and it prevents sudden transitions from $S1$ to $S3$ or vice versa. When the wavelet coefficients fluctuate between values above the higher threshold T_2 and below the lower threshold T_1 in a frequent manner, this indicates the existence of flames in the viewing range of the camera.

In flame pixels, the transition probabilities a should be high and close to each other due to the random nature of uncontrolled fire. On the other hand, transition probabilities should be small in constant-temperature moving bodies, because there is no change or little change in pixel values. Hence we expect a higher probability for b_{00} than any other b value in the nonflame moving pixels model (Fig. 7), which corresponds to higher probability of being in $S1$. The state $S2$ provides hysteresis and it prevents sudden transitions from $S1$ to $S3$ or vice versa.

The transition probabilities between states for a pixel are




Fig. 8 Image frames from some of the test clips. (a), (c), and (d) show fire regions detected and flame boundaries marked with arrows. No false alarms are issued for ordinary moving bright objects in (b), (e), and (f).

estimated during a predetermined period of time around flame boundaries. In this way, the model not only learns the way flame boundaries flicker during a period of time, but also it tailors its parameters to mimic the spatial characteristics of flame regions. The way the model is trained drastically reduces the false alarm rates.

During the recognition phase, the HMM-based analysis is carried out in pixels near the contour boundaries of bright moving regions whose ρ values exceed ρ_T . The state sequence of length of 20 image frames is determined for these candidate pixels and fed to the flame and nonflame pixel models. The model yielding higher probability is determined as the result of the analysis for each of the candidate pixels. A pixel is called as a flame or a nonflame pixel according to the result of this analysis. A fire mask composed of flame pixels is formed as the output of the method.

The probability of a Markov model producing a given sequence of wavelet coefficients is determined by the sequence of state transition probabilities. Therefore, the flame decision process is insensitive to the choice of thresholds T_1 and T_2 , which basically determine if a given wavelet coefficient is close to zero or not.

3 Experimental Results

The proposed method was implemented in a personal computer with an AMD AthlonXP 2000+1.66-GHz processor. The HMMs used in the temporal analysis step were trained using outdoor IR video clips with fire and ordinary moving bright objects like people and cars. Video clips have 236,577 image frames with 160×120 pixel resolution. All of the clips are captured at 10 fps. The FLIR camera that recorded the clips has a spectral range of 8 to 12 μm . Some of the clips were obtained using an ordinary black and



Fig. 9 Image frames from some of the test clips with fire. Pixels on the flame boundaries are successfully detected.

white camera. There are moving cars and walking people in most of the test video clips. Image frames from some of the clips are shown in Figs. 8 and 9.

We used some of our clips for training the Markov models. The fire model was trained with fire videos and the other model was trained with ordinary moving bright objects. The remaining 48 video clips were used for test purposes. Our method yields no false positives in any of the IR test clips.

A modified version of a recent method by Guillemant and Vicente⁶ for real-time identification of smoke in black and white video is implemented for comparison. This method is developed for forest fire detection from watchtowers. In a forest fire, smoke rises first, therefore the method was tuned for smoke detection. Guillemant and Vicente based their method on the observation that the movements of various patterns like smoke plumes produce correlated temporal segments of gray-level pixels, which they called *temporal signatures*. For each pixel inside an envelope of possible smoke regions, they recovered its last d luminance values to form a point $P=[x_0, x_1, \dots, x_{d-1}]$ in d -dimensional "embedding space." Luminance values were quantized to 2^e levels. They utilized fractal indexing using a space-filling Z-curve concept whose fractal rank is defined as:

$$z(P) = \sum_{j=0}^{e-1} \sum_{l=0}^{d-1} 2^{l+jd} x_l^j, \quad (4)$$

where x_l^j is the j 'th bit of x_l for a point P . They defined an instantaneous velocity for each point P using the linked list obtained according to Z-curve fractal ranks. After this step,

Table 1 Detection results for some of the test clips. In the video clip V3, flames are hindered by a wall most of the time.

Video clips	Number of frames with flames	Number of frames in which flames detected		Number of false positive frames	
		CVH method	Our method	CVH method	Our method
V1	0	17	0	17	0
V2	0	0	0	0	0
V3	71	42	63	0	0
V4	86	71	85	0	0
V5	44	30	41	0	0
V6	79	79	79	0	0
V7	0	15	0	15	0
V8	101	86	101	0	0
V9	62	52	59	8	0
V10	725	510	718	54	0
V11	1456	1291	1449	107	0
V12	988	806	981	19	0

they estimated a cumulative velocity histogram (CVH) for each possible smoke region by including the maximum velocity among them, and made smoke decisions about the existence of smoke according to the standard deviation, minimum average energy, and shape and smoothness of these histograms.

Our aim is to detect flames in manufacturing and power plants, large auditoriums, and other large indoor environments. So, we modified the method in Ref. 6 similar to the approach presented in Sec. 2.2. For comparison purposes, we replaced our wavelet-based contour analysis step with

the CVH-based method and left the rest of the algorithm as proposed. We formed two three-state Markov models for flame and nonflame bright moving regions. These models were trained for each possible flame region using wavelet coefficients of CVH standard deviation values. States of HMMs were defined as in Sec. 2.2.

Comparative detection results for some of the test videos are presented in Table 1. The second column lists the number of frames in which flames exist in the viewing range of the camera. The third and fourth columns show the number of frames in which flames were detected by the modified

Table 2 Fire detection results of our method when trained with different flame types.

Video clips	Flame type	Number of frames with flames	Number of frames in which flames detected by our method trained with	
			Paper fire	Paper+alcohol fire
V9	Paper	62	60	58
V10	Paper	725	722	722
V11	Paper+alcohol	1456	1449	1453
V12	Paper+alcohol	993	988	991
V13	Paper+alcohol	1434	1426	1430
V14	Paper	999	996	995

Table 3 Comparison of the proposed method with the modified version of the method in Ref. 6 (CVH method) and the fire detection method described in Ref. 5 for fire detection using a regular visible range camera. The values for processing times per frame are in milliseconds.

Video clips	Number of frames with flames	Number of frames in which flames detected			Processing time per frame (msec)		
		Our method	Method in Ref. 5	CVH method	Our method	Method in Ref. 5	CVH method
V15	37	31	37	26	5	16	11
V16	18	13	18	8	5	17	10
V17	0	2	9	7	4	16	10

CVH method explained in the previous paragraph and our method, respectively. Our method detected flame boundaries that have irregular shapes both temporally and spatially. Both methods detected fire in video clips V3 to V6 and V8 to V12, which contain actual fires indoor and outdoor. In video clip V3, flames were behind a wall most of the time. The distance between the camera and fire ranges between 5 to 50 m in these video clips. Video clips V1, V2, and V7 do not contain any fire. There are flames and people walking in the remaining clips. Some flame frames are missed by both methods, but this is not an important problem at all, because the fire was detected in the next frame or the frame after the next one. The method using CVH detected fire in most of the frames in fire-containing videos as well. However, it yielded false alarms in clips V1, V7, and V9 through V12, in which there were groups of people walking by a car and around a fireplace. The proposed method analyzes the contours of possible fire regions in the wavelet domain. This makes it more robust to slight contour changes than the modified method, which basically depends on the analysis of motion vectors of possible fire regions.

Flames of various burning materials have similar yet different temporal and spatial characteristics. For example, oil flame has a peak flicker frequency around 15 Hz, whereas it is around 5 Hz for coal flame (see Fig. 2 in Ref. 11). In fact, flames of the same burning material under different weather/wind conditions also have different temporal and spatial characteristics. What is common among various flame types is the wide-band random nature in them causing temporal and spatial flicker. Our method exploits this stochastic nature. We use wavelet-based feature signals obtained from flickering flames to train and test our models to get rid of the effects of specific conditions forming the flames. The wavelet domain feature signals capture the condition-independent random nature of flames.

To verify the performance of our method with respect to flames of different materials, we set up the following experiment.

1. We train the model with “paper” fire and test it with both “paper” and “paper+alcohol” fires.
2. We train the model with “paper+alcohol” fire and test it with both “paper” and “paper+alcohol” fires.

Results of this experiment for some of the test clips are

presented in Table 2. The results show that the method has similar detection rates for different fires when trained with different flame types.

The proposed method was also tested with regular video recordings in comparison with the modified version of the method in Ref. 6 and the fire detection method described in Ref. 5. The method in Ref. 5 uses frequency subband analysis to detect 10-Hz flame flicker, instead of using HMMs to capture the random temporal behavior in flames. Results for some of the clips are presented in Table 3. The clip V17 does not contain any fire, either. However, it leads to false alarms because a man with a bright fire-colored shirt dances in front of the camera to fool the algorithm. This man would not cause any false alarms if an infrared camera were used instead of a regular visible range camera.

Notice also that the proposed method is computationally more efficient than Ref. 6 because it is mostly based on contour analysis of the bright moving objects. Average processing time per frame for the proposed method is 5 msec, as shown in Table 3.

4 Conclusion

A novel method to detect flames in IR video is developed. The algorithm uses brightness and motion clues along with a temporal and contour analysis in the wavelet domain. The main contribution of the method is the utilization of hidden Markov models trained using temporal wavelet domain information to detect a random flicker process. The high-frequency behavior of flame region boundaries are analyzed using a wavelet-based contour analysis technique. The experimental results indicate that when the fire falls into the viewing range of an IR camera, the proposed method is successful in detecting the flames without producing false alarms in all the examples that we tried.

The method can be used for both indoor and outdoor early fire detection applications.

Acknowledgments

This work is supported in part by the Scientific and Technical Research Council of Turkey (TÜBİTAK) and European Commission 6th Framework Program with grant number FP6-507752 (MUSCLE Network of Excellence Project). We are grateful to Aselsan Inc. Microelectronics,

Guidance and Electro-Optics Division (MGEO), Ankara, Turkey, for lending "ASIR Thermal Imaging System" and helping us in recording infrared fire videos.

References

1. W. Phillips III, M. Shah, and N. V. Lobo, "Flame recognition in video," *Pattern Recogn. Lett.* **23**(1-3), 319–327 (2002).
2. Fastcom Technology SA, *Method and Device for Detecting Fires Based on Image Analysis*, PCT Pubn.No. WO02/069292, CH-1006 Lausanne, Switzerland (2002).
3. T. Chen, P. Wu, and Y. Chiou, "An early fire-detection method based on image processing," *ICIP '04*, pp. 1707–1710 (2004).
4. B. U. Toreyin, Y. Dedeoglu, and A. E. Cetin, "Flame detection in video using hidden markov models," *ICIP '05*, pp. 1230–1233 (2005).
5. B. U. Toreyin, Y. Dedeoglu, U. Gudukbay, and A. E. Cetin, "Computer vision based system for real-time fire and flame detection," *Pattern Recogn. Lett.* **27**, 49–58 (2006).
6. P. Guillemant and J. Vicente, "Real-time identification of smoke images by clustering motions on a fractal curve with a temporal embedding method," *Opt. Eng.* **40**(4), 554–563 (2001).
7. J. Vicente and P. Guillemant, "An image processing technique for automatically detecting forest fire," *Int. J. Therm. Sci.* **41**, 1113–1120 (2002).
8. B. U. Toreyin, Y. Dedeoglu, and A. E. Cetin, "Wavelet based real-time smoke detection in video," in *13th Eur. Signal Process. Conf. EUSIPCO* (2005).
9. D. S. Chamberlin and A. Rose, "The flicker of luminous flames," in *The First Symp. (Intl.) Combustion, Combustion Institute* (1928).
10. B. W. Albers and A. K. Agrawal, "Schlieren analysis of an oscillating gas-jet diffusion," *Combust. Flame* **119**, 84–94 (1999).
11. J. A. Broadbent, "Fundamental flame flicker monitoring for power plant boilers," *IEE Sem. Adv. Sensors Instrum. Syst. Combustion Process.*, pp. 4/1–4/4 (2000).
12. H. P. Huang, Y. Yan, G. Lu, and A. R. Reed, "On-line flicker measurement of gaseous flames by imaging processing and spectral analysis," *Meas. Sci. Technol.* **10**, 726–733 (1999).
13. G. Healey, D. Slater, T. Lin, B. Drda, and A. D. Goedeke, "A system for real-time fire detection," in *CVPR '93*, pp. 15–17 (1993).
14. R. T. Collins, A. J. Lipton, and T. Kanade, "A system for video surveillance and monitoring," in *8th Intl. Topical Meeting on Robotics and Remote Systems*, American Nuclear Society (1999).
15. Ö. N. Gerek and A. E. Cetin, "Adaptive polyphase subband decomposition structures for image compression," *IEEE Trans. Image Process.* **9**, 1649–1659 (2000).
16. A. E. Cetin and R. Ansari, "Signal recovery from wavelet transform maxima," *IEEE Trans. Signal Process.* **42**, 194–196 (1994).
17. C. W. Kim, R. Ansari, and A. E. Cetin, "A class of linear-phase regular biorthogonal wavelets," *ICASSP '92*, pp. 673–676 (1992).
18. C. B. Liu and N. Ahuja, "Vision based fire detection," *ICPR '04*, **4**, pp. 134–137 (2004).
19. *HMMs Applications in Computer Vision*, H. Bunke and T. Caelli, Eds., World Scientific, New York (2001).



Behcet Uğur Töreyn obtained his MS degree in 2003 from the electrical and electronics engineering department of Bilkent University, Ankara, Turkey, where he is currently a PhD student and a research and teaching assistant. He received his BS degree from the electrical and electronics engineering department of Middle East Technical University, Ankara, Turkey, in 2001. His research interests are in the area of signal processing with an emphasis on image and video processing, pattern recognition, and computer vision. He

has published 19 scientific papers in peer-reviewed conferences and journals.



Ramazan Gökberk Cinbiş is a junior student of computer engineering at Bilkent University, Ankara, Turkey. His research interests include computer vision, pattern recognition, cognitive sciences, and signal processing. He is currently working at RETINA, Vision and Learning Group of Bilkent, where he researches algorithms for content-based image retrieval and mathematical morphology.



Yiğithan Dedeoğlu received BS and MS degrees in computer engineering from Bilkent University, Ankara, Turkey, in 2002 and 2004, respectively. Currently, he is a PhD student in the same department. His research interests are computer vision, pattern recognition, and machine learning.



Ahmet Enis Çetin received the BS degree in electrical engineering from the Middle East Technical University, Ankara, Turkey, and the MSE and PhD degrees in systems engineering from the Moore School of Electrical Engineering, University of Pennsylvania, Philadelphia. From 1987 to 1989, he was an assistant professor of electrical engineering at the University of Toronto, Ontario, Canada. Since then, he has been with Bilkent University, Ankara. Currently, he is a full professor. He is a senior member of the European Association for Signal Processing (EURASIP). Currently, he is a scientific committee member of the EU FP6 funded Network of Excellence (NoE) MUSCLE: Multimedia Understanding through Semantics, Computation and Learning. He is a member of the DSP technical committee of the IEEE Circuits and Systems Society. He founded the Turkish Chapter of the IEEE Signal Processing Society in 1991. He received the Young Scientist Award from the Turkish Scientific and Technical Research Council (TÜBİTAK) in 1993. He is currently on the editorial boards of *EURASIP Journal of Applied Signal Processing* and *Elsevier Signal Processing*.