Integrated Motion Detection and Tracking for Visual Surveillance

Mohamed F. Abdelkader

Rama Chellappa

Qinfen Zheng

Center for Automation Research (CfAR)
University of Maryland,
College Park, MD 20742

Alex L. Chan
U. S. Army Research Laboratory (ARL)
Adelphi, MD

Abstract

Visual surveillance systems have gained a lot of interest in the last few years. In this paper, we present a visual surveillance system that is based on the integration of motion detection and visual tracking to achieve better performance. Motion detection is achieved using an algorithm that combines temporal variance with background modeling methods. The tracking algorithm combines motion and appearance information into an appearance model and uses a particle filter framework for tracking the object in subsequent frames. The systems was tested on a large ground-truthed data set containing hundreds of color and FLIR image sequences. A performance evaluation for the system was performed and the average evaluation results are reported in this paper.

1 Introduction

In the last few years, visual surveillance has become one of the most active research areas in computer vision, especially due to the growing importance of visual surveillance for security purposes. Visual surveillance is a general framework that groups a number of different computer vision tasks aiming to detect, track, and classify objects of interest from image sequences, and on the next level to understand and describe these objects behavior. The ultimate goal in designing smart visual surveillance systems is to replace the existing passive surveillance and to remove, or at least, minimize the need for a human observer to monitor and analyze the visual data.

For these reasons, there has been a number of well-known visual surveillance systems— see [1] for an excellent survey. The real-time visual surveillance system W^4 [2] detects and tracks multiple people and monitors their activi-

ties in an outdoor scene using a single monocular gray-scale or IR camera. It uses a combination of shape analysis and tracking to locate people and their parts, and to create models of their appearance such that they can be tracked through interactions as occlusions. The Pfinder system developed by Wren *et al* [3] builds a 3D description model of a single unoccluded person in a large room using a fixed camera. CMU developed a surveillance system [4] that detects and tracks moving objects over a large area using a network of video sensors; detected objects are classified into categories using shape and color analysis. MIT's system [5] first observes moving objects in a site using a distributed set of sensors, and then uses these observations to learn and classify patterns of activity in the scene that are later used to detect unusual activities.

In this paper, we present an integrated visual surveillance system to monitor an outdoor scene through a stationary forward-looking infra-red (FLIR) or a color camera. Our system begins with the motion detection module, which is responsible for detection and segmentation of moving objects, as well as initializing the tracker. The tracking module estimates the new location of the object at each new frame and maintains the object identity. The performance evaluation module is used to evaluate the system performance using the available ground truth data. The main features of our system are as follows:

- Motion detection algorithm that integrates both temporal variance and background modeling to allow for a robust detection of moving objects.
- An adaptive visual tracking algorithm using both appearance and motion observations in a statistical framework. Particle filters are used in the state estimation to accommodate for appearance changes. The tracking algorithm adaptively updates motion velocity, appearance, noise, and number of particles at each new

frame.

- The integration of the motion detection and tracking stages by using the motion detection module for initialization of the appearance model for the visual tracker, and for obtaining the motion observation.
- Performance evaluation module responsible for measuring and evaluating the performance of the system with respect to a ground- truthed detection data set.

The remainder of the paper is organized as follows. In Section 2, we describe our motion detection algorithm. Section 3 focuses on the process of object tracking using both appearance and motion information. Section 4 presents the performance evaluation of our system. Conclusions are provided in Section 5.

2 Moving-Object Detection

The detection of moving objects is the first stage of a typical surveillance system. Motion detection aims at segmenting the regions pertaining to moving objects from the rest of the image. Subsequent processes such as tracking and behavior analysis are greatly dependent on the performance of this stage.

Many algorithms have been suggested to solve the problem of motion detection, most of which can be categorized into three main approaches: temporal thresholding as in [6], where the moving pixels are identified by thresholding the temporal difference between the frames; background subtraction where detection occurs by comparing the incoming frame with a background model of the scene that is built by modeling the pixel intensity either by a single Gaussian distribution [3], a mixture of Gaussians [5], or using the maximum, minimum and maximum intensity difference as in [2]; and optical flow approaches [7] that use characteristics of flow vectors of moving objects over time to detect moving regions in image sequences.

We present a method based on combining both the temporal variance of the pixel intensities as temporal thresholding approach with background modeling to achieve a robust and accurate motion detection and to reduce false alarms. Our approach is suitable for both color and FLIR sequences collected by stationary surveillance cameras, which are most commonly used in surveillance systems.

2.1 Temporal Variance Based Motion Detection

In our system, we use the temporal variance as a parameter to detect moving areas in a stationary scenes as in [8] and [9]. The mean and variance of the intensity value at each pixel is calculated over a window of several previous

frames and updated recursively for every new frame. This value of the variance is used directly afterward for the detection of moving area. The use of temporal variance for motion detection has some nice properties:

1. The variance of intensity at a certain pixel depends on both the amplitude of changes and the duration of this change so it is more robust to transient noises incurred by moving texture.
2. There is no need for background training period as this method can build the model with the existence of moving objects on the scene.

The mean and variance for the intensity at each pixel (i, j) are recursively updated using a simple exponentially decaying adaptive filter as follows:

$$\begin{aligned} m(i, j, t) &= \alpha m(i, j, t-1) + (1-\alpha)x(i, j, t) \\ m_2(i, j, t) &= \alpha m_2(i, j, t-1) + (1-\alpha)x^2(i, j, t) \\ \sigma^2(i, j, t) &= m_2(i, j, t) - m^2(i, j, t) \end{aligned} \quad (1)$$

where: $x(i, j, t)$ is the intensity, $m(i, j, t)$ is the first moment, $m_2(i, j, t)$ is the second moment and $\sigma^2(i, j, t)$ is the variance at pixel (i, j) at time t , α is the decay rate, that can be rewritten with respect to the filter window size N as:

$$\alpha = \frac{N-1}{N}; N = \frac{1}{1-\alpha} \quad (2)$$

The main problem with using the variance is that it takes a while for the variance to decay back to its original value after the change has ended. This problem causes the moving object to leave a trail behind it, consisting of pixels that were in motion in the proceeding frames. The variance decay rate can be controlled by changing the window size N , however, reducing this size will make the model too adaptive to any changes in the scene. To overcome this problem we propose using a simple background model, which is adaptively updated, to remove this *trail effect*.

2.2 Background modeling

In order to remove the trail effect, a background model is built by recursively updating another set of mean and variance as follows:

$$\begin{aligned} m_{bg}(i, j, t) &= \alpha_{bg}m_{bg}(i, j, t-1) + (1-\alpha_{bg})x(i, j, t) \\ m_{2bg}(i, j, t) &= \alpha_{bg}m_{2bg}(i, j, t-1) + (1-\alpha_{bg})x^2(i, j, t) \\ \sigma_{bg}^2(i, j, t) &= m_{2bg}(i, j, t) - m_{bg}^2(i, j, t) \end{aligned} \quad (3)$$

where: $m_{bg}(i, j, t)$ is the background first moment, $m_{2bg}(i, j, t)$ is the background second moment, and $\sigma_{bg}^2(i, j, t)$ is the background variance for the background

at pixel (i, j) at time t . α_{bg} is the background model decay rate that can also be written with respect to the background filter window size N_{bg} in the same way shown in Equation(2).

The difference between the two models is:

- The window size used for the background model N_{bg} is much larger than that used for the variance update, so that the background model is slowly varying and covers a longer history of the frames.
- The update process for the background model is selective in the sense that only pixels that have not been identified as possible foreground pixels, using the variance, are incorporated at each new frame. So we may assure that this background model does not include any foreground object.

The background model is used to obtain a confidence weight representing the confidence of this pixel being a part of the foreground. This confidence weight is obtained as a function of the distance between the pixel intensity and the background model

$$C(i, j, t) = f \left(\frac{|x(i, j, t) - m_{bg}(i, j, t)|}{\sqrt{\sigma_{bg}(i, j, t)}} \right) \quad (4)$$

where $C(i, j, t)$ is the confidence weight that the pixel (i, j) is part of the foreground and f is a nonlinear-mapping function—such as sigmoid—to map the distance to range of $[0, 1]$ and to emphasize the large distance points. The final binary detection map $L(i, j, t)$ is obtained as follows:

$$L(i, j, t) = \begin{cases} 1 & \text{if } C(i, j, t)\sigma(i, j, t) \geq \text{Threshold} \\ 0 & \text{if } C(i, j, t)\sigma(i, j, t) < \text{Threshold} \end{cases} \quad (5)$$

where the value of Threshold can be obtained either empirically or by multiplying the average background variance by a factor.

2.3 Motion Segmentation

The motion segmentation could be considered as the interface between motion detection and tracking stages of the surveillance system. It includes the segmentation of the moving areas from the binary detection map into disjoint objects, the removal of any small or isolated noise, and the initialization of the bounding boxes that are passed to the subsequent tracker.

In order to perform these tasks, we use the connected-component labeling algorithm presented in [10]. This algorithm performs several raster passes on the binary image $L(t)$ and uses sequential local operations with a one-dimensional table, which produces a fast connected-component result. After the labeling, a bounding box containing all the pixels with the same label is drawn. These

Figure 1. The motion detection stages .(a) The square root of the variance image showing the trails left by the moving objects,(b) The foreground confidence map with the more bright the area the more likely it is a foreground object and (c) The final motion detection map after multiplying (a) and (b) and applying the threshold

bounding boxes are sent to the subsequent tracker in order to start the tracking process.

3 Object tracking

Object tracking module is responsible for estimating the location of each object in each new frame. This module ensures that the object is being tracked even if the motion detection module fails to detect it, either due to occlusion or stopped object.

As a problem of interest for a long time, visual tracking research can be categorized into four main categories [1]; region-based, active-contour-based, feature-based, and model-based tracking. Our tracker belongs to the fourth category, in which we build a model for the tracked object and search for the best match for this model in subsequent frames. Different models have been used, such as: stick figures [11], 2-D contour [12], and appearance models [17]. Different search strategies have been employed to search for the best match, either deterministic by minimizing a cost function as in [13], or probabilistic aiming to estimate the motion state for a state space model using an estimation tool like Kalman filter [14] and Condensation algorithm [15].

In our surveillance system we use a tracking algorithm based on the work presented in [16]. We design a new ob-

servation model that incorporates both appearance and motion information of the object. This algorithm appears to be very effective and robust even in challenging tracking conditions like static occlusion and cluttering background.

The tracking problem can be formulated as an estimation process, where the goal is to estimate the unknown motion state θ_t from a noisy collection of observations, $Y_{1:t} = \{Y_1, \dots, Y_t\}$ arriving in a sequential fashion. For each observed frame Y_t , different image patches Z_t correspond to different motion states θ_t by the relation $Z_t = T\{Y_t; \theta_t\}$, where T is the motion transformation used (Affine in our case). The system transition is usually modeled using a state space model with two important components, the state transition and observation models, which are generally expressed as:

$$\text{state transition model : } \theta_t = F_t(\theta_{t-1}, U_t), \quad (6)$$

$$\text{Observation model : } Y_t = G_t(\theta_t, V_t), \quad (7)$$

where U_t is the system noise, $F_t(\dots)$ characterizes the kinematics, V_t is the observation noise, and $G_t(\dots)$ models the observer. Due to the nonlinear and non-Gaussian nature of our system model; the particle filter [18] is used as a powerful technique to approximate the posterior distribution $p(\theta_t|Y_{1:t})$ by a set of weighted particles $\{\theta_t^{(j)}, w_t^{(j)}\}_{j=1}^J$. Then the state estimate $\hat{\theta}_t$ can be calculated as a maximum a posteriori (MAP) estimate [18].

$$\hat{\theta}_t = \theta_t^{map} = \arg \max_{\theta_t} p(\theta_t|Y_{1:t}) \approx \arg \max_{\theta_t} w_t^{(j)} \quad (8)$$

or any other estimators based on $p(\theta_t|Y_{1:t})$.

3.1 Adaptive State Transition Model

The state in our Model θ_t represent the affine motion parameters used in the system. The state transition model used in our system is a simple adaptive Markov model of the form

$$\theta_t = \hat{\theta}_{t-1} + \nu_t + U_t \quad (9)$$

where U_t is the system noise term and ν_t is the adaptive velocity. We provide a brief overview of the model the readers are referred to [16] for more details.

3.1.1 Adaptive Velocity

Aside from the noise term, the adaptive velocity term $\nu_t \simeq \theta_t - \hat{\theta}_{t-1}$ is predicted using a first-order linear approximation of the constant brightness constraint. A least square solution is found via using the difference between the incoming observation Y_t and the previous set of observations $Z_{t-1} = \{Z_{t-1}^{(j)}\}_{j=1}^J$.

3.1.2 Adaptive Noise and Number of Particles

After calculating the velocity, the variance of the noise U_t and the number of particles in the particle filter are chosen to be proportional to the value of the residual error ϵ_t between the predicted patch $\tilde{Z}_t = T\{Y_t; \hat{\theta}_{t-1} + \nu_t\}$ and the appearance model. This value determines the quality of prediction. Therefore, if ϵ_t is small, implying a good prediction, we only need noise with small variance and small number of particles to absorb the residual motion and vice versa.

3.2 Appearance-Motion Adaptive Observation Model

We use an observation model that combines the motion information (background modeling) and appearance information (foreground modeling) to achieve more robustness especially in the cases where one of the two cues fails. We use the model presented in [19] to combine both cues and assume that the observation Y_t is segmented by the motion state θ_t into two mutually exclusive and collectively exhaustive sets: $\beta(\theta_t)$ and $F(\theta_t)$, which are the regions identified by the state θ_t as background and foreground (target), respectively. With this assumption, the observation equation $p(Y_t|\theta_t)$ can be written as follows:

$$\begin{aligned} p(Y_t|\theta_t) &= p(Y_t(\beta(\theta_t)), Y_t(F(\theta_t))|\theta_t) \\ &= p(Y_t(\beta(\theta_t))|\theta_t)p(Y_t(F(\theta_t))|\theta_t) \quad (10) \\ &= p(Y_t(\beta(\theta_t))|\beta_t)p(Y_t(F(\theta_t))|A_t) \end{aligned}$$

Where A_t is the appearance model and β_t is the background model. Under the assumption that the background is composed of independent pixels.

$$p(Y_t(\beta(\theta_t))|\theta_t)|\beta_t) = \frac{p(Y_t|\beta_t)}{p(Y_t(F(\theta_t))|\beta_t)} \quad (11)$$

So the likelihood function $p(Y_t|\theta_t)$ can be written in the following way

$$p(Y_t|\theta_t) = p(Y_t|\beta_t) \frac{p(Y_t(F(\theta_t))|A_t)}{p(Y_t(F(\theta_t))|\beta_t)} \quad (12)$$

3.2.1 Appearance Modeling

The appearance-based tracking relies on building an adaptive appearance model for the object and estimating the new location of this object in subsequent frames. Our appearance model [16] is an online adaptive appearance model A_t , which is a modified version from the online appearance model (OAM) developed in [17].

Mixture appearance model: we use a time varying OAM that consists of a mixture density of three components $A_t = \{W_t, S_t, F_t\}$. The W -component characterizes a

short time-course, such as in a two-frame tracker, so it is modeled as a Gaussian component that is conditioned on the estimated image patch from the last frame $\hat{Z}_{t-1}$. The fixed template component F is modeled as a Gaussian with fixed mean and variance. Finally, the S -component depicts the stable structure, which is modeled by a slowly varying Gaussian component whose parameters are updated with every new frame to cope for the slow changes in the appearance over time. The observation likelihood is written as

$$p(Y_t(F(\theta_t))|A_t) = p(Z_t|\theta_t) \\ = \prod_{j=1}^d \left\{ \sum_{i=w,s,f} \alpha_{i,t}(j) N(Z_t(j); m_{i,t}(j), \sigma_{i,t}^2(j)) \right\}, \quad (13)$$

where: $\{m_{i,t}; i = w, s, f\}$, $\{\sigma_{i,t}^2; i = w, s, f\}$ and $\{\alpha_{i,t}; i = w, s, f\}$ are the respective mixture means, variances, and mixing probabilities of the three densities, and $N(x; m, \sigma^2)$ is a normal density of mean m and variance σ^2 .

Model Initialization and Update: The appearance model A_1 is initialized by setting $W_1 = S_1 = F_1 = Z_0$, where Z_0 is the the initial area initialized by the detection module.

The update of the appearance model A_t to A_{t+1} is carried out using an Expectation-Maximization (EM) algorithm with the assumption that the past observations are exponentially 'forgotten' with respect to their contribution to the current appearance model.

3.2.2 Motion Modeling

The motion-based modeling in our observation model is formulated by defining a density function $p(Y_t(F(\theta_t))|\beta_t)$ and inserting it into the likelihood equation (12). For this purpose, we use the parameters of the background model introduced in the motion detection algorithm. Assuming that individual pixels are independent,

$$p(Y_t(F(\theta_t))|\beta_t) = p(Z_t|\beta_t) = \prod_{j=1}^d p(Z_t^j|\beta_t) \quad (14)$$

where $p(Z_t^i|\beta_t)$ is the probability distribution of the i^{th} pixel of the patch under the background model. This distribution is a Gaussian $N(Z_t^i; m_{bg}^i, \sigma_{bg}^2)$, where σ_{bg}^2 is a fixed variance indicating the level of confidence in the motion information and determining to what degree this information should contribute to the whole likelihood term.

4 Performance Evaluation

The performance evaluation of surveillance systems has gained more interest in the last few years [20]. Most of the

Table 1. Average performance over the ARL data sequences

Overall TRDR	FAR	Human TDR	Vehicle TDR
82.13%	31.87%	94.09%	83.37%

research in this area is concerned with generating a meaningful error metric to be used in comparing the system results with the ground truth information.

In this section we will provide an evaluation for the results obtained by testing our system on the force protection surveillance system (FPSS) dataset provided by U. S. Army Research Laboratory (ARL). The first set of this data consists of 85 FLIR and color sequences, all of which were manually ground-truthed with respect to the location and type of moving targets on each frame. These sequences represent a wide variety of views, complexity, day times and weather conditions (shine, rain, and snow) of a large parking lot.

In order to evaluate the performance of our surveillance system we adopted some of the performance metrics used in [21] such as : Tracker Detection Rate (TRDR), False Alarm Rate (FAR), and Track Detection Rate (TDR)

$$\begin{aligned} TRDR &= \frac{TotalTruePositives}{TotalNumberOfGroundTruthPoints} \\ FAR &= \frac{TotalFalsePositives}{TotalTruePositives + TotalFalsePositives} \\ TDR &= \frac{Numberoftruepositivesfortrackedobject}{Totalnumberofgroundtruthpointsforobject} \end{aligned} \quad (15)$$

Where a **True Positive** is defined as a ground truth point that is located within the bounding box of an object detected and tracked by the tracking algorithm. A **False negative** is a ground truth point that is not located with the bounding box of any object tracked by the tracking algorithm. A **False positive** is an object that is tracked by the system that does not have a matching ground truth point.

(**TRDR**) and (**FAR**) characterize the overall tracking performance of the motion detection and object tracking algorithms for all the objects, while, **TDR** indicates the tracking completeness of a specific ground truth track. Because of the space limitation, we will just present the average evaluation result for all the sequences (Table.1).

5 Conclusion

We have described a visual surveillance system that integrates motion detection and tracking tasks to achieve a better tracking performance. Temporal variance of pixel intensities is used for motion detection combined with

Figure 2. Example for different object being tracked over time (Red boxes) along with the ground truth data where green, blue and yellow crosses represent human, vehicle, and other objects, respectively

background modeling to remove trail effect. We used a visual tracking algorithm that combines motion and appearance information into an observation model, uses an adaptive state transition model with adaptive velocity and noise, and uses particle filter as a tool for estimating the unknown state vector. The system was tested on a large data set of FLIR and color sequences collected by ARL FPSS program and the performance was evaluated with respect to manually ground-truthed data.

Acknowledgement

This work was supported in part by the Advanced Sensors Consortium sponsored by the U.S. Army Research Laboratory under the Collaborative Technology Alliance Program, Cooperative Agreement DAAD19-01-2-0008

References

- [1] W. Hu, T. Tan, L. Wang, and S. Maybank. *A survey on visual surveillance of object motion and behaviors*, IEEE Trans. on systems, man, and cybernetics part C: Applications and Reviews, Vol.34, NO.3, AUGUST 2004.
- [2] I. Haritaoglu, D. Harwood, and L. S. Davis. *W⁴: Real-time surveillance of people and their activities*, IEEE Trans. PAMI, vol. 22, pp. 809-830, Aug. 2000.
- [3] C. R. Wren, A. Azarbayejani, T. Darrell, and A. P. Pentland. *Pfinder: real-time tracking of the human body*, IEEE Trans. PAMI., vol. 19, pp. 780-785, July 1997.
- [4] R. T. Collins, A. J. Lipton, T. Kanade, H. Fujiyoshi, D. Duggins, Y. Tsin, D. Tolliver, N. Enomoto, O. Hasegawa, P. Burt, and L. Wixson. *A system for video surveillance and monitoring*, Carnegie Mellon Univ., Pittsburgh, PA, Tech. Rep., CMU-RI-TR-00-12, 2000.
- [5] W. E. L. Grimson, C. Stauffer, R. Romano, and L. Lee. *Using adaptive tracking to classify and monitor activities in a site*, in Proc. IEEE Conf. CVPR, Santa Barbara, CA, 1998, pp. 22-31.
- [6] A. J. Lipton, H. Fujiyoshi, and R. S. Patil. *Moving target classification and tracking from real-time video*, in Proc. of Fourth IEEE Workshop on Applications of Computer Vision, 1998. WACV '98
- [7] J. Barron, D. Fleet, and S. Beauchemin. *Performance of optical flow techniques*, International Journal of Computer Vision, 12(1):42-77, 1994.
- [8] Q. Zheng and S. Der. *Moving target indication in LRAS3 sequences*, in 5th Annual Fedlab Symposium College Park MD, 2001
- [9] S. Joo and Q. Zheng. *A Temporal Variance-Based Moving Target Detector*, in Proceedings of the IEEE International Workshop on Performance Evaluation of Tracking and Surveillance (PETS), Jan. 2005
- [10] K. Suzuki, I. Horiba, and N. Sugie. *Fast Connected-Component Labeling Based on Sequential Local Operations in the Course of Forward Raster Scan Followed by Backward Raster Scan*, IEEE International Conference on Pattern Recognition (ICPR'00)- Volume 2, p. 24-34
- [11] I. A. Karaulova, P. M. Hall, and A. D. Marshall. *A hierarchical model of dynamics for tracking people with a single video camera*, in Proc. British Machine Vision Conf., 2000, pp. 262-352.
- [12] S. Ju, M. Black, and Y. Yacobi. *Cardboard people: a parameterized model of articulated image motion*, in Proc. IEEE Int. Conf. Automatic Face and Gesture Recognition, 1996, pp. 38-44.
- [13] G. D. Hager, and P. N. Belhumeur. *Efficient region tracking with parametric models of geometry and illumination*, IEEE Trans. PAMI, vol. 20, pp. 1025-1039, Oct. 1998.
- [14] T. J. Brodia, S. Chandra, and R. Chellappa. *Recursive techniques for estimation of 3-d translation and rotation parameters from noisy image sequences*, IEEE Trans. Aerosp. Electron. Syst., vol. 26, pp. 639-656, Apr. 1990.
- [15] M. Isard and A. Blake. *Contour tracking by stochastic propagation of conditional density*, in Proc. Eur. Conf. Computer Vision, 1996.
- [16] S. Zhou, R. Chellappa, and B. Moghaddam. *Visual tracking and recognition using appearance-adaptive*

- models in particle filters*. IEEE Transactions on Image Processing (IP), Vol. 11, pp. 1434-1456, November 2004.
- [17] A. D. Jepson, D. J. Fleet, and T. El-Maraghi. *Robust online appearance model for visual tracking*, in Proc. IEEE Conf. CVPR, vol. 1, 2001, pp. 415-422.
 - [18] A. Doucet, N. D. Freitas, and N. Gordon. *Sequential Monte Carlo Methods in Practice*. New York: Springer-Verlag, 2001.
 - [19] Aswin C. Sankaranarayanan, Rama Chellappa and Qinfen Zheng. *Tracking Objects in Video Using Motion and Appearance Models*, To appear in IEEE International Conference on Image Processing (ICIP 2005).
 - [20] T.J. Ellis. *Performance Metrics and Methods for Tracking in Surveillance*, Proceedings of the Third International Workshop on Performance Evaluation of Tracking and Surveillance (PETS2002), Copenhagen, June 2002.
 - [21] J. Black, T. J. Ellis, and P. Rosin. *A novel method for video tracking performance evaluation*. In Joint IEEE Int. Workshop on Visual Surveillance and Performance Evaluation of Tracking and Surveillance (VS-PETS), pages 125-132, October 2003.