



A Fast Model-Free Morphology-Based Object Tracking Algorithm

Jonathan Owens^a, Andrew Hunter^b & Eric Fletcher^a

^aSchool of Computing & Technology

University of Sunderland, UK

jonathan.owens@sunderland.ac.uk

^bDepartment of Computer Science

Durham University, UK

andrew1.hunter@durham.ac.uk

Abstract

This paper describes the multiple object tracking component of an automated CCTV surveillance system. The system tracks objects, and alerts the operator if unusual trajectories are discovered. Objects are detected by background differencing. Low contrast levels can present problems, leading to poor object segmentation and fragmentation, particularly on older analogue surveillance networks. The model-free tracking algorithm described in this paper addresses object fragmentation, and the object merging that occurs when proximate objects segment to the same connected component.

1 Introduction

Automated visual surveillance aims to provide an attention-focussing filter to enable an operator to make an optimum decision whenever an unusual event occurs [1]. This is achieved by directing the operator's attention only to those events classified as unusual. The backbone of such systems typically comprises something like the processing pipeline shown in figure 1.




Figure 1: Typical image processing pipeline for automated video surveillance

The blocks outlined in bold are dealt with in this paper, focussing on the object tracking module, which must deal with the uncertainty of object segmentation. This uncertainty is manifest when moving objects are segmented by background differencing, where it is common for the segmented object to fragment due to parts of the object matching the greyscale of the background. This problem is exacerbated when CCTV system managers wish to implement modern automated surveillance techniques



on top of the existing surveillance infrastructure. Older cameras are typically low resolution, monochrome, analogue devices with CCD arrays of low dynamic range, producing images of low contrast.

Even complex multimodal background representations cannot successfully segment objects if these closely match the background. Typical object tracking algorithms employ the Kalman filter or some other *predictor-corrector* iterative algorithm for dealing with uncertainty in the tracking plane [2], [3], [4], [5], [6]. When the uncertainty of object segmentation is too great, the tracking algorithm relies on the *predictor* step, and the state of the tracked object is updated using the internal model, rather than the observed measurements. It is an implicit assumption in typical background differencing and tracking algorithms that segmentation is successful at maintaining objects holistically, without fragmentation.

When objects merge in the binary difference image, *predictor-corrector* algorithms may use partial image evidence to update an object state, but rely more heavily on the internal model of object parameters, and do so until such time as the objects separate and can be tracked individually. This problem can be overcome, together with other phenomena such as occlusion, with an explicit model fit to tracked objects [5], [7], [8].

In this paper, a simple, fast object tracking algorithm is described which attempts to maintain the morphology of tracked objects, given the evidence provided by the segmentation block of the pipeline. The algorithm is intended for use with sparsely occupied scenes with low activity levels. This restriction is not prohibitive to the objective of the system, as unusual activity (e.g. car crime) typically occurs in otherwise quiescent scenes. This algorithm is part of a hybrid novelty detection system [9], [10]. The overall philosophy of the system is that it should be self-organising, requiring no user defined models of scene elements, object forms or object motion. A self-organising map is used to measure the novelty of a vector describing local motion, while a hierarchical network classifies the global pattern of object motion.

2 Summary of Algorithm

The method of choice for moving object segmentation in most tracking algorithms is background differencing. Methods based on the calculation of optic flow are computationally intensive and use a raw image feature match to maintain a track if the object stops moving [11], [12]. The multiple gaussian per pixel representation used in [2] is robust, but entails a huge computational cost. To keep the background generator computable in real-time on non-specialised hardware, we employ a simple low-pass filter [13], [14].

The CCTV images are obtained from an analogue, monochrome camera at a resolution of 640x480 pixels, with a colour range of 256 grey values. The low-pass filter method is able to cope with slow changes in luminance, such as the movement of shadows cast by static objects. At time t , the difference and background images are calculated as follows:

$$\delta(r, c) = \begin{cases} 1 & |I(r, c) - B(r, c)| > T \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

$$B(t) = \beta \times I(t) + (1 - \beta) \times B(t-1) \quad \text{if } \delta(r, c) = 0 \quad (2)$$

where r and c are the row and column subscripts of a single pixel, δ is the difference image, I is the input image, and B is the background image. If the difference $\delta(r, c)$ is

greater than the threshold, $T = 12$, the pixel is labelled as foreground, while background pixels are modified by expression 2, where $\beta=0.1$ and controls the rate of the background update. Sufficiently rapid luminance changes cause background differencing to fail. To guard against this eventuality, the system counts the total number of pixels assigned to the foreground and if this is greater than 50% of the image, the entire background image is reset and tracking is restarted.

Noise is removed by applying a morphological “opening” operator to the difference image. The pixels that remain classified as foreground are collected into 4-connected components and assigned unique identities. Examples of pedestrians and their segmented silhouettes are shown in figure 2. Along with an identity, each object has an associated feature vector, the elements of which are *area*, *width*, *height* and a histogram of the greyscale distribution of object pixels. This feature vector is used to match objects from frame to frame, as described in detail in the following section. (In this paper, the term ‘silhouette’ describes a single connected component, of which there are three instances in each of the two images in the central column in figure 2).




Figure 2: Examples of partially segmented silhouettes and the objects they delineate

3 Object Tracking

The object tracker described here is a purely measurement based object-to-silhouette matching algorithm with morphological manipulation that deals with uncertainty in the segmentation algorithm. The philosophy of the tracking algorithm is motivated by general top-down assumptions of the types of unusual behaviour and normal activity that the system will have to deal with.

Based on the surveillance of a typical car park scene, the following assumptions are made: (1) Activity of pedestrians is of primary interest; vehicles are tracked but their activity will not be passed to the “behaviour” classifier modules. (2) Vehicle crime, the main form of novel behaviour that is of interest given the monitored scene, is most likely to be carried out by independent pedestrians. (3) Pedestrians entering the scene in a group, i.e. with very similar temporal and spatial origins with respect to the tracking plane, are likely to have a common origin and destination. Hence, tracking the centroid of the group will give a reasonable approximation to individual trajectories within the group. (4) Pedestrians with differing spatial and temporal origins may also have differing intended destinations. A distinct history for each object should be maintained, even if the silhouettes of such objects merge. Based on the above assumptions, the tracker attempts to track objects in the form in which they are initially segmented. Therefore, the system will try to maintain the tracking of distinct objects, even if their

silhouettes merge with those of other objects. This entails the use of a silhouette-partitioning function to separate merged silhouettes prior to the best-match process.

The algorithm will also try to maintain the overall morphology of a group, which entails the use of a silhouette-combiner which attempts to maintain the track of a group as a single entity. This function serves a two-fold purpose. If a group of pedestrians is being tracked, a global track on the whole group will be maintained by merging any group members who temporarily separate from the group silhouette. On the other hand, if the segmentation of an object fails and it becomes fragmented into distinct components, the silhouette-combiner will attempt to gather the fragments together until the best match to the object is achieved (see figure 2).

After the background differencing step, the binary image consists of silhouettes, where a number of silhouettes may correspond to a single object (fragmentation), or a single silhouette may “cover” more than one object (merging). Associated with each object and silhouette is a feature vector, $\mathbf{f}_i = [a, w, h, \mathbf{g}]$ where \mathbf{g} is a 16 element vector of the greyscale histogram covered by silhouette i (figure 3c,d), a is the area, or number of pixels making up the silhouette and w and h are the width and height of the minimum bounding rectangle (figure 3a).




Figure 3: (a) Object delineated by minimum bounding rectangle calculated from binary silhouette (b). (d) Greyscale histogram of segmented object (c)

Central to the algorithm described below is the concept of “difference” between object and silhouette feature vectors. The difference between the feature vectors of object Q and silhouette S is defined as

$$\mathbf{d}(Q, S) = \left[|a_Q - a_S|, |h_Q - h_S|, |w_Q - w_S|, \sqrt{(\mathbf{g}_Q - \mathbf{g}_S) \bullet (\mathbf{g}_Q - \mathbf{g}_S)} \right] \quad (3)$$

which is a four element vector comprised of the absolute differences of the area, height and width of the object and silhouette and the scalar length of the greyscale histogram difference vector. When tracking begins (or after the scene has been empty), there will be no objects to which the newly segmented silhouettes can be compared; in this case, the algorithm will jump to *step 8*, where sufficiently large silhouettes that are not matched to tracked objects instantiate new entries in the object list. Otherwise, the algorithm proceeds by trying to match silhouettes conservatively to existing objects, by preferentially matching silhouettes to established objects, then to new objects, and then improving the match by resolving silhouette merging and fragmentation.

Step 1 – Naïve Match: The distances in the image plane between the centroid of every object, \mathbf{q} , and the centroids of the segmented silhouettes, \mathbf{s} , are used to form a valid-match matrix, V , based on an arbitrary search radius around position \mathbf{q} . The search radius, $r = 80$ pixels, establishes a limit on the number of possible matches that can be



evaluated by the object-to-silhouette assignment algorithm. V is a matrix of dimension $\{n, m\}$, where n is the number of tracked objects and m is the number of segmented silhouettes. Thus, V is constructed as follows

$$V_{ji} = 1 \text{ if } \|\mathbf{q}_j, \mathbf{s}_i\| \leq r \quad (1 \leq j \leq n, 1 \leq i \leq m) \quad (4)$$

where \mathbf{q} and \mathbf{s} are the vectors describing the position in the image plane of the object j and silhouette i . The “cost” of every object-to-silhouette assignment having a non-zero entry in matrix V is given by the scalar value

$$c(Q, S) = \sum_k \frac{\mathbf{d}_k(Q, S)}{\mathbf{f}_k(Q)} \quad (5)$$

where $\mathbf{d}_k(Q, S)$ is the k^{th} element of the difference vector between object Q and silhouette S , and $\mathbf{f}_k(Q)$ is the feature vector of object Q . The histogram element in the feature vector in the denominator of expression (5) is transformed into a scalar value by calculating the Euclidean length of the vector. The object feature vector $\mathbf{f}_k(Q)$, scales the elements of the difference vector, assuming that the within population coefficients of variation are roughly equal for the separate feature vector elements.

The match matrix is initialised by assigning silhouettes to objects, on a per-object basis, where a single silhouette may be matched to more than one object. It should be noted that, in expression 6, an object-to-silhouette assignment is contingent upon a corresponding non-zero entry in valid-match matrix V :

$$M_{jv} = V_{jv} \text{ where } v = \arg \min_i \left\{ c(Q_j, S_i) \right\} \quad (1 \leq j \leq n, 1 \leq i \leq m) \quad (6)$$

where a non-zero entry indicates a match between object Q_j and silhouette S_i . Ideally, the naïve match would be enough to unambiguously match objects to the segmented silhouettes. The remaining steps of the algorithm are designed to address the errors that may arise from silhouette fragmentation and merging.

Step 2 – Remove Duplicate Matches: As the naïve match is allocated by choosing the lowest cost match per object, there exists the potential for match conflicts, where a silhouette is initially matched with more than one object. At this stage, objects are allocated to one of two classes, transient or non-transient. Transient objects are those that have only been instantiated for one frame – an object must find a silhouette match over two frames before it is classified as non-transient. If there are match conflicts, silhouette matches to transient object are removed if these overlap with matches to non-transient objects. This step makes it less likely that false objects will interfere with the tracking of real objects. False objects may be attributed to noise or interaction of the object with the environment, such as reflections on vehicles.

Step 3 – Evaluate Possible Merges: Where match conflicts arise between non-transient objects, the duplicate match may be caused by the merging of silhouettes in close proximity. To establish whether a merging event has occurred, the objects with conflicting silhouette assignments are combined into a single macro-object, Θ , with a single feature vector. If the cost of the macro-object to silhouette match is lower than the minimum single object to silhouette match cost, then it is assumed that the silhouettes have merged. The cost of the macro-object to silhouette match is given by

$$c(\Theta, S) = \sum_k \frac{\mathbf{d}_k(\Theta, S)}{\mathbf{f}_k(Q_j)} \text{ where } j = \arg \min_v \left\{ c(Q_v, S_i) \right\} \quad (7)$$

The cost of the match between macro-object Θ and silhouette S is the summation of the elements in the difference vector $\mathbf{d}(\Theta, S)$ scaled by the corresponding elements in the

feature vector $\mathbf{f}(Q_j)$. Q_j is the individual object that best matches the silhouette, the cost of this match being $c(Q_j, S)$, calculated in step 1. If $c(\Theta, S) < c(Q_j, S)$, it is assumed that the silhouettes of the objects have merged and the conflict is left to be resolved in step 5. Otherwise, the match to object Q_j is retained and the other objects are assigned to their next lowest cost match. The reassignment of objects could raise further match conflicts between non-transient objects so the algorithm repeats steps 2 and 3 until conflicts are removed or found to result from silhouette merging.

Step 4 – Remove Duplicate Matches: Non-transient object matches modified in the last part of step 3 may have been allocated to silhouettes matched to transient objects. This is permitted because established objects take priority over transient objects as it is possible these may simply be a product of a patch of noise in the last frame. Transient object matches that conflict with the relocated non-transient objects are removed. Each silhouette allocated to a non-transient object is labelled as “securely matched”, and the difference vectors, $\mathbf{d}(Q, S)$, recalculated ready for the next step.

Step 5 – Partition Merged Silhouettes: Here we apply the first of the morphological refinement algorithms; the silhouette-partitioning function is applied to resolve silhouettes matched to more than one non-transient object. If a duplicated match got past step 3, it is likely that the silhouettes of the objects have merged and require separating to allow the tracker to maintain a separate track of each object.

Silhouette Partition Function: Given the (x, y) co-ordinates of each pixel in the silhouette, the sum of least squares linear regression line, $y = a + bx$ can be calculated directly from the following expressions:

$$a = \frac{(\sum y)(\sum x^2) - (\sum x)(\sum xy)}{n_p \sum x^2 - (\sum x)^2}, \quad b = \frac{n_p \sum xy - (\sum x)(\sum y)}{n_p \sum x^2 - (\sum x)^2} \quad (8)$$

where the expressions are applied to all n_p pixels in the silhouette.

Each pixel is projected onto the linear regression line giving a histogram of the silhouette’s distribution of mass along the line. The silhouette is divided by placing partition lines at intervals along the regression line, l_r . To calculate the partition points, the sizes of the objects (at time $t-1$) participating in the split are listed in according to their relative positions along the x-axis. Given n objects, there will be $n-1$ partitions \mathbf{p} , based on the distribution of “mass” among the n objects. If the left-most extent of the silhouette projection on to the regression line is the origin, and the right-most extent is unity, the partitions p_m will lie in the range $\{0, 1\}$,

$$p_m = \frac{\sum_{j=1}^m a_j}{\sum_{k=1}^n a_k} \quad (1 \leq m \leq n-1) \quad (9)$$

where p_m is partition m , a_j is the area of object j and n is the total number of objects participating in the split. Moving from left to right along the silhouette regression line histogram, the p_m ratios are used to place the partition points relative to the total mass of the merged silhouette. Each pixel belonging to the silhouette is labelled according to its projection onto the regression line (figure 4a). Calculating the partition intervals with expression (9) assumes that the overlap between the participating objects is not significant. Large overlaps will mean the partitions are offset with an error that increases as we progress from left to right along l_r , as illustrated in figure 4b.




Figure 4: (a) The linear regression line (white) and the resulting partition shown by the grey levels mapping the partitioned silhouettes. (b) Heavy mutual occlusion may distort partitioned silhouettes (right-most figure)

Based on the pixel labels, a feature vector is calculated for each sub-silhouette. The differences between the feature vectors of the participating objects and the new partitioned silhouettes are calculated and the costs evaluated with the expression

$$c(Q, \Phi) = \sum_k \frac{\mathbf{d}_k(Q, \Phi)}{\mathbf{f}_k(Q^{AV})} \quad (10)$$

where $\mathbf{d}(Q, \Phi)$ is the difference between object Q and a partitioned sub-silhouette Φ , and the cost of the match is the summation of the scaled elements of the difference vector. The object Q^{AV} whose feature vector is used to scale the elements of the difference vectors is simply an average object calculated from the feature vectors of the objects participating in the partition of the silhouette. Among the n objects and n sub-silhouettes involved in the silhouette split, the new object to silhouette matches are assigned on a greedy lowest-cost basis, with match conflicts obviously not permitted at this stage.

The object match matrix M is adjusted to accommodate the new silhouettes and revised match assignments, and these can take part in the subsequent steps in exactly the same way as unmodified silhouettes.

Step 6 – Merge Fragmented Silhouettes: Here, the possibility of object fragmentation is addressed, in which an object may appear as several separate silhouettes in the binary difference image. Silhouettes matched to non-transient objects are combined with silhouettes lacking a secure match (as defined in *step 4*) within the valid search radius. By combining fragments into a merged silhouette Ψ , a fragmented silhouette may be reconstructed, improving the match to the tracked object. The cost of a new match is evaluation with the expression

$$c(Q, \Psi) = \sum_k \frac{\mathbf{d}_k(Q, \Psi)}{\mathbf{f}_k(Q)} \quad (11)$$

If $c(Q, \Psi) < c(Q, S)$, that is if the merged silhouette Ψ produces a closer match to object Q than the unmodified silhouette S , the merge is accepted and the silhouettes are combined into a single entry in the silhouette list. This process continues until all unallocated silhouette fragments within the valid match radius have been considered.

Step 7 – Refine Transient Matches: The so-called transient objects may have been instantiated over a patch of noise in the previous frame, or they may be genuine new objects entering the field of view. A cost-reducing feature combination step is performed across these objects as in *step 6*, i.e. at this stage only transient objects are examined and may only be combined with unallocated silhouettes. This priority given to persistent objects is one way to reduce the susceptibility of the overall system to short-lived noise and temporary object fragmentation.



Step 8 – Update Objects: Given the match matrix M , the object lists are updated. Objects without a match are removed and each unassigned silhouette of sufficient size instantiates a new object in the list. The size criterion helps to prevent persistent noise, which is usually comprised of small image patches, from instantiating an object.

5 Object Based Reference Update

The stationarity of non-pedestrian objects is determined to assist in maintaining a valid reference image. When an object is stationary for >16 frames (i.e. 4 seconds at the 4Hz sampling interval) the object is inserted into the background image, pedestrians typically sway even when standing, so inserted objects are typically parked vehicles. The previously determined minimum bounding rectangle of the silhouette is used to define the region of the input that is copied to the background.

Once an object has been inserted into the background, its object list entry is transferred to a “recently-inserted-object” list, and foreground objects, i.e. pedestrians, can now be tracked as they pass in front of or exit the vehicle. If the event was a “drop-off”, rather than a parking event, the vehicle will subsequently move away from its previously stationary position, leaving a “hole” in the background. The negative object will be detected as being stationary, and the centroid can be compared to those in the “recently-inserted-object” list. If the distance between the stationary false object centroid and a list entry is below a threshold, the object is inserted immediately into the background, thereby patching the “hole” in the reference image as quickly as possible.

This stationary object reference update is useful because of the assumption that the system will only submit pedestrian activity to the novelty detection components, thereby dictating that tracking localises pedestrians at the expense of tracking other objects.

6 Performance of the Object Tracker

The object tracking algorithm was evaluated with respect to the monitored scene as interpreted by a human observer, the overall description of which could be called the ‘operator perceived activity’ (OPA). The operator looked for discrepancies between actual activity and that “perceived” by the tracker. The system was evaluated on 3 days of live video from 8:00am to 10:00am, comprising a total of 6 hours, spanning a range of activity levels, from peak activity to relatively quiescent periods. The tracker performance is shown in table 1. The left side of the table summarises results for pedestrian events, and the right shows the vehicle events. From a total of 311 separate events, 264 were tracked perfectly, i.e. an accuracy of 84.9% . Figure 5 shows examples of fragmented and merged objects disambiguated by the tracker. Instances in which there was a discrepancy between the tracker and the OPA are discussed below.

Correctly tracked events lie down the main diagonal of both sections of the table; e.g. there were 120 instances where a single pedestrian was correctly tracked. The 27 entries where one pedestrian was present but two pedestrians were tracked refers to the situation where a pedestrian fragmented and one segment was momentarily tracked as a separate pedestrian – this situation was temporary and the extra transient track did not interfere with the tracking of the true object. The three instances of complete tracking failure (cell $\{0,1\}$ of table 1 (a)) resulted from excessively poor segmentation producing fragments too small to instantiate an object list entry.

		(a) Number of Pedestrians				(b) Number of Vehicles				
		Tracker				Tracker				
		0	1	1	3		0	1	2	3
OPA	0		9	1	0	0		0	0	0
	1	3	120	27	2	1	0	139	1	0
	2	0	3	4	0	2	0	0	0	0
	3	0	0	1	1	3	0	0	0	0

Table 1: Performance of the object tracker, comparing the number of vehicles and pedestrians tracked (columns) with the actual events as defined by an operator (rows)



Figure 5: Fragmented (top) and merged (bottom) objects successfully disambiguated

The 9 instances of pedestrians being tracked when in fact there were none (cell $\{0,1\}$ of table 1 (a)) was due temporary regions generated by phenomena such as reflections of pedestrians on vehicle windows, detached shadows from vehicles or elongated fragments of vehicles. The single instance where a vehicle was incorrectly tracked (cell $\{1,2\}$ of table 1 (b)) was due to a neatly fragmented vehicle giving rise to two vehicle sized objects that were tracked separately. It should be noted that pedestrian activity is not submitted to the novelty detection networks until the pedestrian has been tracked coherently for approximately 3 seconds, so the entries in the left side of table 1 lying above the main diagonal, showing tracking false pedestrians, were not passed on to the novelty detection modules. From the point of view of activity classification, significant tracking errors were those lying below the main diagonal in the left side of table 1. These were instances where the tracker “lost” the track on one or more pedestrians, thereby rendering them “invisible” to the novelty detectors. Therefore, considering only those table entries lying on or below the main diagonal, out of 132 separate pedestrian events, 125 were successfully passed to the classifier stages of the surveillance system, giving a tracking accuracy of 94.6%.

7 Discussion

The tracking algorithm attempts to maintain objects in the form in which they were instantiated, which is achieved by means of two morphological operators – A merging operator deals with silhouette fragmentation and a partitioning function handles



silhouette merging events. As shown in table 1, sometimes the fragmentation of objects is so poor that a perfect track cannot be maintained. However, this is dealt with by the next module in the processing pipeline. Indeed, by accepting the motion data only from objects that have been in existence for a given period, the novelty detection modules [9], [10], can prevent false transient objects from generating alarms.

The algorithm is able to track poorly segmented objects on the basis of form only, and no prior models of size, shape or texture are used. This is consistent with the overall strategy of a self-organising system, where objects are tracked and their behaviour is classified without *a priori* knowledge built into the system. The algorithm is extremely fast, as the elements used during the match process are simple macroscopic features of silhouette area, width, height and greyscale histogram.

8 References

- [1] Foresti, G.L., Mähönen, P., Regazzoni, C.S. (eds): Multimedia Video-Based Surveillance Systems – Requirements, Issues and Solutions. Kluwer Academic Publishers
- [2] Stauffer, C., Grimson, W.E.L.: Learning Patterns of Activity Using Real-Time Tracking. IEEE Trans. PAMI, Vol. 22, No. 8 (2000)
- [3] Foresti, G.L.: A Real-Time System for Video Surveillance of Unattended Outdoor Environments. IEEE Trans. Circuits and Systems for Vid Tech, Vol. 8, No. 6 (1998)
- [4] Foresti, G.L., Roli, F.: Learning and Classification of Suspicious Events for Advanced Visual-Based Surveillance. In: Foresti, G.L., Mähönen, P., Regazzoni, C.S. (eds): Multimedia Video-Based Surveillance Systems: Requirements. Kluwer Academic Publishers
- [5] Remagnino, P., Baumberg, A., Grove, T., Hogg, D., Tan, T., Worrall, A., Baker, K.: An Integrated Traffic and Pedestrian Model-Based Vision System. Proc. BMVC (1997)
- [6] Rasmussen, C., Hager, G.D.: Probabilistic Data Association Methods for Tracking Complex Visual Objects. IEEE Trans. PAMI, Vol. 23, No. 6 (2001)
- [7] Ferryman, J.M., Worrall, A.D., Sullivan, G.D., Baker, K.D.: Visual Surveillance Using Deformable Models of Vehicles. Robotics and Autonomous Systems, Vol. 19, No. 3-4 (1997)
- [8] Pece, A., Worrall, A.: A Statistically-based Newton Method for Pose Refinement. Image and Vision Computing, Vol. 16, No. 8 (1998)
- [9] Owens, J., Hunter, A.: Application of the Self-Organising Map to Trajectory Classification. IEEE Third International Workshop on Visual Surveillance (2000)
- [10] Owens, J., Hunter, A., Fletcher, E.: Novelty Detection in Video Surveillance Using Hierarchical Neural Networks. Proc. ICANN (to appear) (2002)
- [11] Boghossian, B.A., Velastin, S.A.: Image Processing System for Pedestrian Monitoring Using Neural Classification of Normal Motion Patterns. Measurement and Control, Vol. 32, No. 9 (1999)
- [12] Medioni, G., Cohen, I., Bremond, F., Hongeng, S., Nevatia, R.: Event Detection and Analysis from Video Streams. IEEE Trans. PAMI, Vol. 23, No. 8 (2001)
- [13] Makarov, A.: Comparison of Background Extraction Based Intrusion Detection Algorithms. IEEE Int. Conf. Image Processing (1996)
- [14] Kehtarnavaz, N., Rajkotwala, F.: Real-Time Vision-Based Detection of Waiting Pedestrians. Real-Time Imaging, Vol. 3 (1997)