

Research Article

Active Video Surveillance Based on Stereo and Infrared Imaging

Gabriele Pieri and Davide Moroni

Institute of Information Science and Technologies, Via G. Moruzzi 1, 56124 Pisa, Italy

Correspondence should be addressed to Gabriele Pieri, gabriele.pieri@isti.cnr.it

Received 28 February 2007; Accepted 22 September 2007

Recommended by Eric Pauwels

Video surveillance is a very actual and critical issue at the present time. Within this topics, we address the problem of firstly identifying moving people in a scene through motion detection techniques, and subsequently categorising them in order to identify humans for tracking their movements. The use of stereo cameras, coupled with infrared vision, allows to apply this technique to images acquired through different and variable conditions, and allows an a priori filtering based on the characteristics of such images to give evidence to objects emitting a higher radiance (i.e., higher temperature).

Copyright © 2008 G. Pieri and D. Moroni. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

1. INTRODUCTION

Recognizing and tracking moving people in video sequences is generally a very challenging task, and automatic tools to identify and follow a human “target” are often subject to constraints regarding the environment under investigation, the characteristics of the target itself, and its full visibility with respect to the background.

Current approaches regarding real-time target tracking are based on (i) successive frame differences [1], using also adaptive threshold techniques [2], (ii) trajectory tracking, using weak perspective and optical flow [3], and (iii) region approaches, using active contours of the target and neural networks for movement analysis [4], or motion detection and successive regions segmentation [5]. In recent years, thanks to the improvement of infrared (IR) technology and the drop of its cost, also thermal infrared imagery has been widely used in tracking applications [6, 7]. Besides, the fusion of visible and infrared imagery is starting to be explored as a way to improve the tracking performance [8].

Regarding specific approaches for human tracking, frame difference, local density maxima, and human shape models are used in [9, 10] for tracking in crowded scenes, while face and head tracking by means of appearance-based methods and background subtraction are used in [11].

For the surveillance of wide areas, there is a need of multiple-cameras coordination, in [12], there is a posterior integration of the different single cameras tracks in a global track using a probabilistic multiple-camera model.

In this paper, the problem of detecting a moving target and its tracking is faced by processing multisource information acquired using a vision system capable of stereo and IR vision. Combining the two acquisition modalities assures different advantages consisting, first of all, of an improvement of target-detection capability and robustness, guaranteed by the strength of both media as complementary vision modalities. Infrared vision is a fundamental aid when low-lighting conditions occur or the target has similar colour to the background. Moreover, as a detection of the thermal radiation of the target, the IR information can be manageably acquired on a 24-hour basis, under suitable conditions. On the other hand, the visible imagery, when available, has a higher resolution and can supply more detailed information about target geometry and localization with respect to the background.

The acquired multisource information is firstly elaborated for detecting and extracting the target in the current frame of the video sequence. Then the tracking task is carried on using two different computational approaches. A hierarchical artificial neural network (HANN) is used during active tracking for the recognition of the actual target, while,

when the target is lost or occluded, a content-based retrieval (CBR) paradigm is applied on an a priori defined database to relocalize the correct target.

In the following sections, we describe our approach, demonstrating its effectiveness in a real case study, the surveillance of known scenes for unauthorized access control [13, 14].

2. PROBLEM FORMULATION

We face the problem of tracking a moving target distinguishable from a surrounding environment owing to a difference of temperature. In particular, we consider overcoming lighting and environmental condition variation using IR sensors.

Humans tracking in a video sequence consists of two correlated phases: target spatial localization, for individuating the target in the current frame, and target recognition, for determining whether the identified target is the one to be followed.

Spatial localization can be subdivided into detection and characterization, while recognition is performed for an active tracking of the target, frame by frame, or for relocalizing it, by means of an automatic target search procedure.

The initialization step is performed using an automatic motion-detection procedure. A moving target appearing in the scene under investigation is detected and localized using the IR camera characteristics, and eventually the visible cameras under the hypothesis to be working in a known environment with known background geometry. A threshold, depending on the movement area (expressed as the number of connected pixels) and on the number of frames in which the movement is detected, is used to avoid false alarms. Then the identified target is extracted from the scene by a rough segmentation. Furthermore, a frame-difference-based algorithm is used to extract a more detailed (even if more subject to noise) shape of the target.

Once segmented, the target is described through a set of meaningful multimodal features, belonging to morphological, geometric, and thermographic classes computed to obtain useful information on shape and thermal properties.

To cope with the uncertainty of the localization, increased by partial occlusions or masking, an HANN can be designed to process the set of features during an active tracking procedure in order to recognize the correctness of the detected target.

In case the HANN does not recognize the target, wrong object recognition should happen due to either a masking, partial occlusion of the person in the scene, or a quick movement in an unexpected direction. In this circumstance, the localization of the target is performed by an automatic search, supported by the CBR on a reference database. This automatic process is considered only for a dynamically computed number of frames, and, if problems arise, an alert is sent and the control is given back to the user.

The general algorithm implementing the above-described approach is shown in Figure 1 and it regards its on-line processing. In this case, the system is used in real time to perform the tracking task. Extracted features from the selected target drive active tracking with HANN and support




FIGURE 1: Automatic tracking algorithm.

the CBR to resolve the queries to the database in case of lost target. Before this stage, an off-line phase is necessary, where known and selected examples are presented to the system so that the neural network can be trained, and all the extracted multimodal features can be stored in the database, which is organised using predefined semantic classes as the key. For each defined target class, sets of possible variations of the initial shape are also recorded, for taking into account that the target could be still partially masked or have a different orientation. More details of the algorithm are described as follows.

3. TARGET SPATIAL LOCALIZATION

3.1. Target detection

After the tracking procedure is started, a target is localized and segmented using the automatic motion-detection procedure, and a reference point, called centroid C_0 , internal to it is selected (e.g., the center of mass of the segmented object detected as motion can be used for the first step). This point is used in the successive steps, during the automatic detection, to represent the target. In particular, starting from C_0 , a motion-prediction algorithm has been defined to localize the target centroid in each frame of the video sequence. According to previous movements of the target, the current expected position is individuated, and then refined through a neighborhood search, performed on the basis of temperature-similarity criteria.

Let us consider the IR image sequence $\{F_i\}_{i=0,1,2,\dots}$, corresponding to the set of frames of a video, where $F_i(p)$ is the thermal value associated to the pixel p in the i th frame. The trajectory followed by the target, till the i th frame, $i > 0$, can

```

Function Prediction ( $i, \{F_i\}_{i=0,1,2,\dots}, n$ );
/* Check if the target has moved over a threshold
distance in last  $n$  frames
if  $\|C_{i-n} - C_{i-1}\| > Threshold_1$ 
then
/* Compute the expected target position  $P_i^1$ 
in the current frame by interpolating the last  $n$ 
centroid positions
 $P_i^1 = \text{INTERPOLATE}(\{C_j\}_{j=i-n,\dots,i-1})$ ;
/* Compute the average length of the movements
of the centroid
 $d = (\sum_{j=i-n}^{i-2} \|C_j - C_{j+1}\|) / (n - 1)$ ;
/* Compute a new point on the basis of temperature
similarity criteria in a circular neighborhood
 $\Theta_d$  of  $P_i^1$  of radius  $d$ 
 $P_i^2 = \arg \min_{P \in \Theta_d} [F_i(P) - F_{i-1}(C_{i-1})]$ ;
if  $\|P_i^1 - P_i^2\| > Threshold_2$ 
then
 $P_i^3 = \alpha P_i^1 + \beta P_i^2$ ; /* where  $\alpha + \beta = 1$ 
/* Compute the final point in acircular
neighborhood  $N_r$  of  $P_i^3$  of radius  $r$ 
 $C_i = \arg \min_{P \in N_i} [F_i(P) - F_{i-1}(P_{i-1})]$ ;
else
 $c_i = P_i^2$ ;
else /* Compute the new centroid according to
temperature similarity in a circular
neighborhood  $N_1$  of the last centroid
 $C_i = \arg \min_{P \in N_i} [F_i(P) - F_{i-1}(P_{i-1})]$ 
Return  $C_i$ 

```

ALGORITHM 1: Prediction algorithm used to compute the candidate centroid in a frame.

be represented as the centroids succession $\{C_j\}_{j=0,\dots,i-1}$. The prediction algorithm for determining the centroid C_i in the current frame can be described as shown in Algorithm 1.

Where i is the sequential number of the current frame, $\{F_i\}$ is the sequence of frames, the number of frames considered for prediction is the last n , and $F_i(P)$ represents the temperature of point P in the i th frame.

The coordinates of centroids referring to the last n frames are interpolated for detecting the expected position P_i^1 . Then, in a circular neighborhood of P_i^1 of radius equal to the average movement amplitude, an additional point P_i^2 is detected as the point having the maximum similarity with the centroid C_{i-1} of the previous frame. If $\|P_i^2 - P_i^1\| > Threshold_2$, then a new point P_i^3 is calculated as a linear combination of the previous determined ones. Finally, a local maximum search is again performed in the neighborhood of P_i^3 to make sure that it is internal to a valid object. This search finds the point C_i that has the thermal level closest to the one of C_{i-1} .

Starting from the current centroid C_i , an automated edge segmentation of the target is performed using a gradient descent along 16 directions starting from C_i . Figure 2 shows a sketch of the segmentation procedure and an example of its result.

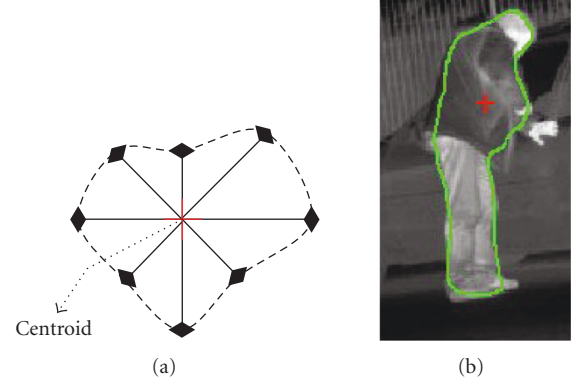


FIGURE 2: Example of gradient descent procedure to segment a target (a) and its application to an example frame identifying a person (b).

3.2. Target characterization

Once the target has been segmented, multisource information is extracted in order to obtain a target description. This is made through a feature-extraction process performed on the three different images available for each frame in the sequence. The sequence of images is composed of both grey-level images (i.e., frames or thermographs) of a high-temperature target (with respect to the rest of the scene) integrated with grey-level images obtained through a reconstruction process [15].

In particular, the extraction of a depth index from the grey-level stereo images, performed by computing disparity of the corresponding stereo points [16], is realized in order to have significant information about the target spatial localization in the 3D scene and the target movement along depth direction, which is useful for the determination of a possible static or dynamic occlusion of the target itself in the observed scene.

Other features, consisting in radiometric parameters measuring the temperature and visual features, are extracted from the IR images. There are four different groups of visual features which are extracted from the region enclosed by the target contour defined by the sequence of N_c (i.e., in our case, $N_c = 16$) points having coordinates $\langle x_i, y_i \rangle$.

Semantic class

The semantic class the target belongs to (i.e., an upstanding, crouched, or crawling person) can be considered as an additional feature and is automatically selected, considering combinations of the above-defined features, among a predefined set of possible choices and assigned to the target.

Moreover, a class-change event is defined, which is associated with the target when its semantic class changes in time (different frames). This event is defined as a couple $\langle SC_b, SC_a \rangle$ that is associated with the target, and represents the modification from the semantic class SC_b selected before and the semantic class SC_a selected after the actual frame, important features to consider in order to retrieve when the semantic class of the target changes are the morphological

features, and in particular, an index of the normal histogram distribution.

Morphological: shape contour descriptors

The morphological features are derived extracting characterization parameters from the shape obtained through frames difference during the segmentation.

To avoid inconsistencies and problems due to intersections, the difference is made over a temporal window of three frames.

Let $\Delta(i-1, i)$ be the modulus of difference between the frames F_{i-1} and F_i . Otsu's thresholding is applied to $\Delta(i-1, i)$ in order to obtain a binary image $B(i-1, i)$. Letting TS_i to be the target shape in the frame F_i , heuristically we have

$$B(i-1, i) = TS_{i-1} \cup TS_i. \quad (1)$$

Thus the target shape is approximated for the frame at time i by the formula

$$TS_i = B(i-1, i) \cap B(i, i+1). \quad (2)$$

Once the target shape is extracted, first, an edge detection is performed in order to obtain a shape contour, and second, a computation of the normal in selected points of the contour is performed in order to get a better characterization of the target. These steps are shown in Figure 3.

Two morphological features, the normal orientation and the normal curvature degree, based on the work by Berretti et al. [17], are computed. Considering the extracted contour, 64 equidistant points $\langle s_i, t_i \rangle$ are selected. Each point is characterized by the orientation θ_i of its normal and its curvature K_i . To define these local features, a local chart is used to represent the curve as the graph of a degree 2 polynomial. More precisely, assuming without loss of generality that, in a neighborhood of $\langle s_i, t_i \rangle$, the abscissas are monotone, the fitting problem

$$t = as^2 + bs + c \quad (3)$$

is solved in the least square sense. Then we define

$$\theta_i = a \tan\left(\frac{-1}{2as_i + b}\right), \quad (4)$$

$$K_i = \frac{2a}{(1 + (2as_i + b)^2)^{3/2}}.$$

Moreover, the histogram of the normal orientation, discretized into 16 different bins, corresponding to the same directions above mentioned is extracted.

Such a histogram, which is invariant for scale transformation and thus independent of the distance of the target, will be used for a deeper characterization of the semantic class of the target. This distribution represents an additional feature to the classification of the target, for example, a standing person will have a far different normal distribution than




FIGURE 3: Shape extraction by frames difference (top), edge detection superimposed on the original frame (centre), and boundary with normal vector on 64 points (bottom). Left and right represent two different postures of a tracked person.

a crawling one (see Figure 4), a vector $[v(\theta_i)]$ of the normal for all the points in the contour is defined, associated to a particular distribution of the histogram data.

Geometric

$$\text{Area} = \frac{\left| \sum_{i=1}^{N_c} [(x_i y_{i+1}) - (y_i x_{i+1})] \right|}{2}, \quad (5)$$

$$\text{Perimeter} = \sum_{i=1}^{N_c} \sqrt{(x_i - x_{i+1})^2 + (y_i - y_{i+1})^2}.$$

Thermographic

$$\text{Average Temp: } \mu = \frac{1}{\text{Area}} \sum_{p \in \text{Target}} F_i(p),$$

$$\text{Standard dev.: } \sigma = \sqrt{\frac{1}{\text{Area} - 1} \sum_{p \in \text{Target}} (F_i(p) - \mu)^2},$$

$$\text{Skewness: } \gamma_1 = \frac{\mu_3}{\mu_2^{3/2}},$$

$$\text{Kurtosis: } \beta_2 = \frac{\mu_4}{\mu_2^2},$$

$$\text{Entropy: } E = - \sum_{p \in \text{Target}} F_i(p) \log_2(F_i(x, y)),$$

(6)

where μ_r are moments of order r .

All the extracted information is passed to the recognition phase in order to assess if the localized target is correct.

3.3. Target recognition

The target recognition procedure is realised using a hierarchical architecture of neural networks. In particular, the architecture is composed of two independent network levels, each using a specific network typology that can be trained separately.

The first level focuses on clustering the different features extracted from the segmented target; the second level performs the final recognition, on the basis of the results of the previous one

The clustering level is composed of a set of classifiers, each corresponding to one of the aforementioned classes of features. These classifiers are based on unsupervised self-organizing maps (SOM) and the training is performed to cluster the input features into classes representative of the possible target semantic classes. At the end of the training, each network is able to classify the values of the specific feature set. The output of the clustering level is an m -dimensional vector consisting of the concatenation of the m SOMs outputs (in our case, $m = 3$). This vector represents the input of the second level.

The recognition level consists of a neural network classifier based on error backpropagation (EBP). Once trained, such network is able to recognize the semantic class that can be associated to the examined target. If the semantic class is correct, as specified by the user, the detected target is recognized and the procedure goes on with the active tracking. Otherwise, wrong target recognition occurs and the automatic target search is applied to the successive frame in order to find the correct target.

3.4. Automatic target search

When wrong target recognition occurs, due to masking, occlusion, or quick movements in unexpected directions, the automatic target search starts.

The multimodal features of the candidate target are compared to the ones recorded in a reference database. A simi-

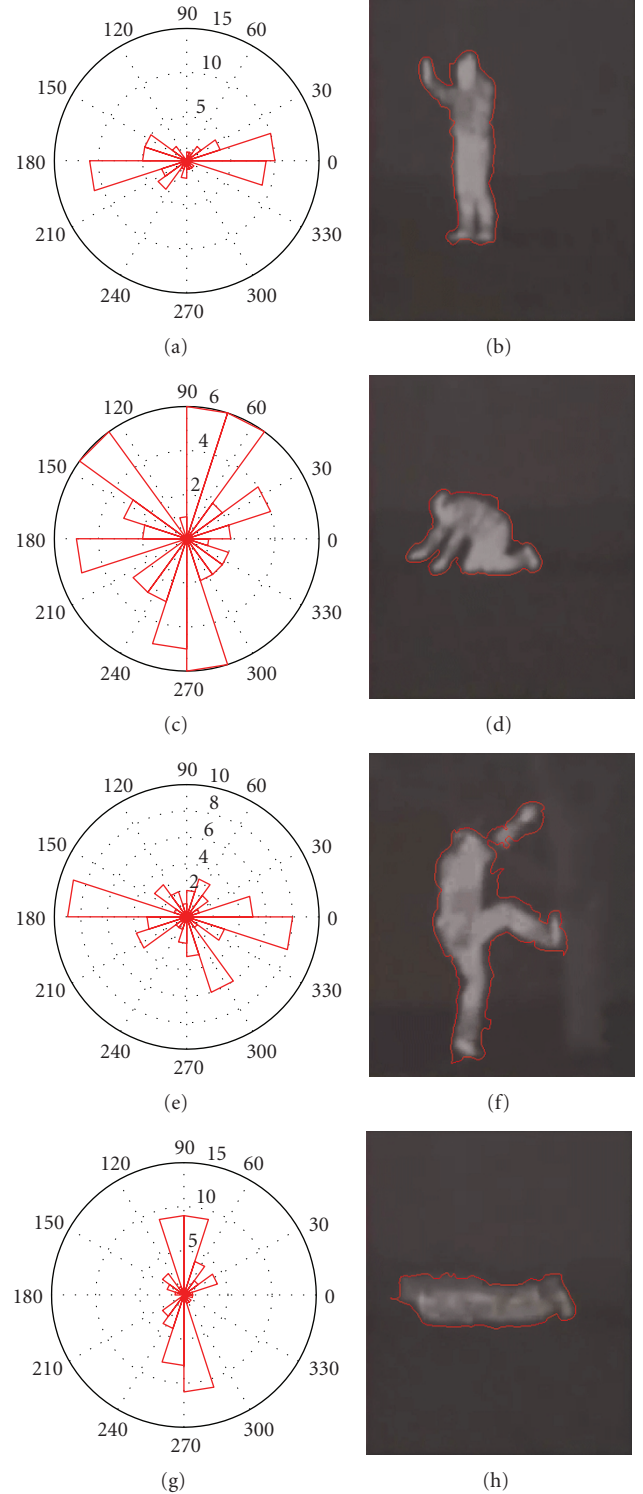


FIGURE 4: Distribution histogram of the normal (left) of targets having different postures (right).

ilarity function is applied for each feature class [18]. In particular, we considered colour matching, using percentages and colour values, and shape matching, using the cross-correlation criterion, and the vector $[v(\theta_i)]$ representing the distribution histogram of the normal.

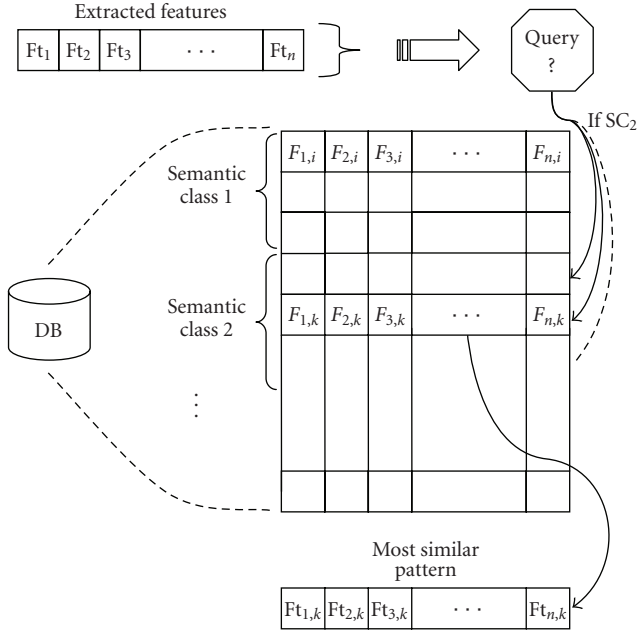


FIGURE 5: Automatic target search supported by a reference database and driven by the semantic class feature to restrict the number of records.

In order to obtain a global similarity measure, each similarity percentage is associated to a preselected weight, using the reference semantic class as a filter to access the database information.

For each semantic class, possible variations of the initial shape are recorded. In particular, the shapes to compare with are retrieved in the MM database using information in a set obtained considering the shape information stored at the time of the initial target selection joined with the one of the last valid shape.

If the candidate target shape has a distance, from at least one in the obtained set, below a fixed tolerance threshold, then it can be considered valid. Otherwise, the search starts again in the next frame acquired [13].

In Figure 5, a sketch of the CBR, in case of automatic target search, is shown considering with the assumption that the database was previously defined (i.e., off-line), and considering a comprehensive vector of features $\{F_{t_k}\}$ for all the above-mentioned categories.

Furthermore, the information related to a semantic class change is used as a weight for possible candidate targets; this is done considering that a transition from a semantic class SC_b to another class SC_a has a specific meaning (e.g., a person who was standing before and is crouched in the next frames) in the context of a surveillance task, which is different from other class changes.

The features of the candidate target are extracted from a new candidate centroid, which is computed starting from the last valid one (C_v). From C_v , considering the trajectory of the target, the same algorithm as in the target-detection step is applied so that a candidate centroid C_i in the current frame is found and a candidate target is segmented.




FIGURE 6: Tracking of a target person moving and changing posture (from left to right: standing, crouched, and crawling).

With respect to the actual feature vector, if the most similar pattern found in the database has a similarity degree higher than a prefixed threshold, then the automatic search has success and the target tracking for the next frame is performed through the active tracking. Otherwise, in the next frame, the automatic search is performed again, still considering the last valid centroid C_v as a starting point.

If, after j_{MAX} frames, the correct target has not yet been grabbed, the control is given back to the user. The value of j_{MAX} is computed considering the Euclidean distance between C_v and the edge point of the frame E_r along the search direction r , divided by the average speed of the target previously measured in the last f frames $\{C_j\}_{j=0,\dots,v}$ (7),

$$j_{MAX} = \frac{\|C_v - E_r\|}{\left(\sum_{j=v-f}^{v-1} \|C_j - C_{j+1}\|/f\right)}. \quad (7)$$

4. RESULTS

The method implemented has been applied to a real case study for video surveillance to control unauthorized access in restricted-access areas.

Due to the nature of the targets to which the tracking has been applied, using IR technology is fundamental. The temperature that characterizes humans has been exploited to enhance the contrast of significant targets with respect to a surrounding background.

The videos were acquired using a thermo camera in the 8–12 μm wavelength range, mounted on a moving structure covering 360° pan and 90° tilt, and equipped with 12° and 24° optics to have 320 × 240 pixel spatial resolution.

Both the thermo-camera and the two stereo high-resolution visible cameras were positioned in order to explore a scene 100-meter far, sufficient in our experimental environments. The frame acquisition rate ranged from 5 to 15 fps.

In the video-surveillance experimental case, during the off-line stage, the database was built taking into account different image sequences relative to different classes of the monitored scenes. In particular, the human class has been composed taking into account three different postures (i.e., upstanding, crouched, and crawling) considering three



FIGURE 7: Example of an identified and segmented person during video surveillance on a gate.



FIGURE 8: Example of an identified and segmented person during video surveillance in a parking lot.

different people typologies (short, middle, and tall) (see Figure 6).

A set of surveillance videos were taken during night time and positioned in specific areas, such as a closed parking lot and an access gate to a restricted area, for testing the efficiency of the algorithms. Both areas were under suitable illumination conditions to exploit visible imagery.

The estimated number of operations, performed for each frame when tracking persons, consists of about $5 \cdot 10^5$ operations for the identification and characterization phases, while the active tracking requires about $4 \cdot 10^3$ operations. This assures the real-time functioning of the procedure on a personal computer of medium power. The automatic search process can require a higher number of operations, but it is performed when the target is partially occluded or lost due to some obstacles, so it can be reasonable to spend more time in finding it, thus losing some frames. Of course, the number of operations depends on the relative dimension of the target to be followed, that is, bigger targets require a higher effort to be segmented and characterized.

Examples of persons tracking and class identification are shown in Figures 7 and 8.

The acquired images are preprocessed to reduce the noise.

5. CONCLUSION

A methodology has been proposed for detection and tracking of moving people in real-time video sequences acquired with two stereo visible cameras and an IR camera mounted on a robotized system.

Target recognition during active tracking has been performed, using a hierarchical artificial neural network (HANN). The HANN system has a modular architecture which allows the introduction of new sets of features including new information useful for a more accurate recognition. The introduction of new features does not influence the training of the other SOM classifiers and only requires small changes in the recognition level. The modular architecture allows the reduction of local complexity and, at the same time, the implementation of a flexible system.

In case of automatic searching of a masked or occluded target, a content-based retrieval paradigm has been used for the retrieval and comparison of the currently extracted features with the previously stored in a reference database.

The achieved results are promising for further improvements as the introduction of additional new characterizing features and enhancement of hardware requirements for a quick response to rapid movements of the targets.

ACKNOWLEDGMENTS

This work was partially supported by the European Project Network of Excellence MUSCLE—FP6-507752 (Multimedia Understanding through Semantics, Computation and Learning). We would like to thank M. Benvenuti, head of the R&D Department at TD Group S.p.A., for his support and for allowing the use of proprietary instrumentation for test purposes. We would also like to thank the anonymous referee for his/her very useful comments.

REFERENCES

- [1] A. Fernandez-Caballero, J. Mira, M. A. Fernandez, and A. E. Delgado, "On motion detection through a multi-layer neural network architecture," *Neural Networks*, vol. 16, no. 2, pp. 205–222, 2003.
- [2] S. Fejes and L. S. Davis, "Detection of independent motion using directional motion estimation," *Computer Vision and Image Understanding*, vol. 74, no. 2, pp. 101–120, 1999.
- [3] W. G. Yau, L.-C. Fu, and D. Liu, "Robust real-time 3D trajectory tracking algorithms for visual tracking using weak perspective projection," in *Proceedings of the American Control Conference (ACC '01)*, vol. 6, pp. 4632–4637, Arlington, Va, USA, June 2001.
- [4] K. Tabb, N. Davey, R. Adams, and S. George, "The recognition and analysis of animate objects using neural networks and active contour models," *Neurocomputing*, vol. 43, pp. 145–172, 2002.
- [5] J. B. Kim and H. J. Kim, "Efficient region-based motion segmentation for a video monitoring system," *Pattern Recognition Letters*, vol. 24, no. 1–3, pp. 113–128, 2003.
- [6] M. Yasuno, N. Yasuda, and M. Aoki, "Pedestrian detection and tracking in far infrared images," in *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern*

- Recognition (CVPR '04)*, pp. 125–131, Washington, DC, USA, June–July 2004.
- [7] J. Zhou and J. Hoang, “Real time robust human detection and tracking system,” in *Proceedings of the 2nd Joint IEEE International Workshop on Object Tracking and Classification in and Beyond the Visible Spectrum*, San Diego, Calif, USA, June 2005.
- [8] B. Bhanu and X. Zou, “Moving humans detection based on multi-modal sensory fusion,” in *Proceedings of IEEE Workshop on Object Tracking and Classification Beyond the Visible Spectrum (OTCBVS '04)*, pp. 101–108, Washington, DC, USA, July 2004.
- [9] C. Beleznai, B. Fruhstuck, and H. Bischof, “Human tracking by mode seeking,” in *Proceedings of the 4th International Symposium on Image and Signal Processing and Analysis (ISPA '05)*, vol. 2005, pp. 1–6, Nanjing, China, November 2005.
- [10] T. Zhao and R. Nevatia, “Tracking multiple humans in complex situations,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 26, no. 9, pp. 1208–1221, 2004.
- [11] A. Utsumi and N. Tetsutani, “Human tracking using multiple-camera-based head appearance modeling,” in *Proceedings of the 6th IEEE International Conference on Automatic Face and Gesture Recognition (AFGR '04)*, pp. 657–662, Seoul, Korea, May 2004.
- [12] T. Zhao, M. Aggarwal, R. Kumar, and H. Sawhney, “Real-time wide area multi-camera stereo tracking,” in *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR '05)*, vol. 1, pp. 976–983, San Diego, Calif, USA, June 2005.
- [13] M. G. Di Bono, G. Pieri, and O. Salvetti, “Multimedia target tracking through feature detection and database retrieval,” in *Proceedings of the 22nd International Conference on Machine Learning (ICML '05)*, pp. 19–22, Bonn, Germany, August 2005.
- [14] S. Colantonio, M. G. Di Bono, G. Pieri, O. Salvetti, and M. Benvenuti, “Object tracking in a stereo and infrared vision system,” *Infrared Physics and Technology*, vol. 49, no. 3, pp. 266–271, January 2007.
- [15] M. Sohail, A. Gilgiti, and T. Rahman, “Ultrasonic and stereo vision data fusion,” in *Proceedings of the 8th International Multitopic Conference (INMIC '04)*, pp. 357–361, Lahore, Pakistan, December 2004.
- [16] O. Faugeras and Q.-T. Luong, *The Geometry of Multiple Images*, The MIT press, Cambridge, Mass, USA, 2004.
- [17] S. Berretti, A. Del Bimbo, and P. Pala, “Retrieval by shape similarity with perceptual distance and effective indexing,” *IEEE Transactions on Multimedia*, vol. 2, no. 4, pp. 225–239, 2000.
- [18] P. Tzouveli, G. Andreou, G. Tsechpenakis, Y. Avrithis, and S. Kollias, “Intelligent visual descriptor extraction from video sequences,” in *Proceedings of the 1st International Workshop on Adaptive Multimedia Retrieval (AMR '04)*, vol. 3094 of *Lecture Notes in Computer Science*, pp. 132–146, Hamburg, Germany, September 2004.