An Integrated Multi-Modal Sensor Network for Video Surveillance

Andrea Prati, Roberto Vezzani
University of Modena and Reggio Emilia
{prati.andrea, vezzani.roberto}@unimore.it

Luca Benini, Elisabetta Farella, Piero Zappi
University of Bologna
{lbenini, efarella, pzappi}@deis.unibo.it

ABSTRACT

To enhance video surveillance systems, multi-modal sensor integration can be a successful strategy. In this work, a computer vision system able to detect and track people from multiple cameras is integrated with a wireless sensor network mounting PIR (Passive InfraRed) sensors. The two subsystems are briefly described and possible cases in which computer vision algorithms are likely to fail are discussed. Then, simple but reliable outputs from the PIR sensor nodes are exploited to improve the accuracy of the vision system. In particular, two case studies are reported: the first uses the presence detection of PIR sensors to disambiguate between an opened door and a moving person, while the second handles motion direction changes during occlusions. Preliminary results are reported and demonstrate the usefulness of the integration of the two subsystems.

Categories and Subject Descriptors

I.4.8 [**Computing Methodologies**]: Image Processing and Computer Vision—*Scene Analysis*

General Terms

Design,Algorithms,Security

Keywords

Sensor Network, PIR, multiple cameras, tracking, video surveillance

1. INTRODUCTION

Video surveillance and other security-related applications have gained many credits due to the terroristic threats of the last years. Several industrial and academic projects have recently started to increase the accuracy of (semi-)automatic surveillance systems. In addition, the abatement of hardware costs allows the deployment of thousands of cameras for surveillance purposes at a reasonable cost.

The ever-increasing demand of security and the low cost of cameras contributed to the diffusion of the research in *distributed multi-camera surveillance systems*. Multiple cameras enable the surveillance of wider areas and the exploitation of redundant information (provided by the different viewpoints) might solve classical limitations of single-camera systems, such as occlusions.

Despite the efforts made by the researchers in developing a robust multi-camera vision system, computer vision algorithms have proven their limits to work in complex and cluttered environments. These limits are mainly due to two classes of problems. The first is that "non-visible areas can not be processed by the system". This trivial statement is of particular importance in cluttered scenes and can be partially lessened by using multiple sensors (not only cameras). The second class of problems, instead, is due to the limited resolution of cameras. Having infinite resolution and zooming capabilities would make the job easier, but, in addition to be unfeasible, it would exponentially increase the computational load and it is typically too expensive.

An interesting solution to these problems is that of using simple but effective specialized sensors to solve the specific problems of the vision systems. In this way, vision would still provide high-level information, and low-level sensors would assure higher accuracy. In this context, the marriage between a widely distributed low-cost wireless sensor network and the coarsely distributed higher level of intelligence that can be exploited by computer vision systems may overcome many troubles in a complete tracking of large areas. For our application, we exploit Passive Infrared (PIR) sensors which are widely deployed in low-cost surveillance systems (e.g., for car or home alarm systems). PIR sensors are used in traditional surveillance systems to trigger the activation of video cameras [9]. A trigger just conveys a binary (yes/no) presence information, but limited signal processing effort on the output of a PIR sensor can produce much more information (e.g., target speed and direction of movement). Furthermore, integration of data from multiple networked PIR sensors can provide drastically improved spatial resolution in monitoring and tracking. Low-cost and low-power sensor nodes can now be developed with on-board processing capabilities, reconfigurability and wireless connectivity [5, 10].

This work reports our joint research in developing a multimodal sensor network that integrates a wireless network of PIR-based sensors with a traditional vision system to provide drastically improved (in accuracy and robustness) tracking of people. It is worth noting that people surveil-

VSSN'05, November 11, 2005, Singapore.

lance is more interesting from the research's point of view w.r.t. to vehicle tracking, because of the intrinsic complexity in detecting, tracking, and understanding human behaviour: changes in posture and gestures, human interaction, presence of multiple people, and so on, make the problem challenging and interesting for the computer vision community.

The paper is structured as follows. The next section analyzes the related works on the topics of distributed video surveillance, multi-modal sensor integration and PIR sensors presented in the literature. Sections 3 and 4 describe the two subsystems of our prototype, namely the PIR sensor network and the computer vision subsystem, respectively. Their integration and preliminary results are provided in section 5. Eventually, conclusions are drawn.

2. RELATED WORKS

The emerging technology of multisensor data fusion has a wide range of applications, both in public and in private areas. Among them, surveillance and security have gained much interest in recent years, also due to the terroristic threats. Such applications often rely on cameras, due to the large amount of information and the high-level knowledge they provide. However, video surveillance in wide area through computer vision techniques is still a challenging problem largely faced in the last years [7]. Additionally to the problems of segmentation and tracking of a single camera video surveillance system, it requires a consistent labeling phase. In other words, the identities of the objects detected from each camera must be preserved in order to analyze their behaviours.

Adjacent cameras can have overlapped field of views or not. In the case of non-overlapped cameras the only feature utilizable for maintaining the identity of a moving subject is its visual appearance. To this aim color and shape ratios are widely adopted [18]. Unfortunately, using merely the subject's appearance is not a successful strategy, since the appearance (in particular, the colour) can be reproduced very differently with different cameras and under different illumination conditions. Some works that use only subject's appearance have been proposed also in the case of overlapped cameras [13, 14]. However, in the case of partially overlapped cameras, the best choice is to exploit also geometrical information. This class of solutions are called geometry-based approaches and they can be further subdivided into calibrated and uncalibrated. In [16], each camera processes the scene and obtains a set of tracks. Then, regions along the epipolar lines in pairs of cameras are matched and the mid-points of the matched segments are back-projected in 3D and then, with an homography, onto the ground plane to identify possible positions of the person within a probability distribution map (filtered with a Gaussian kernel). A particularly interesting paper is reported in [21] in which homography is exploited to solve occlusions. Single camera processing is based on particle filter and on probabilistic tracking based on appearance to detect occlusions.

Most of these approaches require camera calibration. In outdoor environments with many cameras, placed in high positions over poles at unknown distance, manual calibration could be difficult and time consuming to achieve. Thus, automatic camera calibration techniques have been proposed. Relevant example of these are the work of Jaynes [11] and the one of Khan and Shah [12]. The first one supposes that the trajectory of the people during the transition between a camera's field of view to that of another camera keeps the same direction and the same speed. With this assumption the paper shows a method to register the two views. The approach of Khan and Shah, instead, is based on the computation of the so-called *Edges of Field of View* (EOFOV, hereinafter), i.e. the lines delimiting the field of view of each camera and, thus, defining the overlapped regions. Through a learning procedure in which a single track moves from one view to another, an automatic procedure computes these edges that are exploited to keep consistent labels on the objects when they pass from one camera to the adjacent.

Our algorithm for consistent labelling follows the same geometrical approach of this proposal to compute, starting from the EOFOV lines extraction, the homographic relationship between the two ground planes in an automatic way. This transformation is then exploited for establishing the consistent labelling.

Since cameras still have limitations due to the limited field of view or resolution, the proposal of this work is to combine data coming from cameras with information provided by a wireless sensor network based on Passive Infrared Sensors (PIR). PIR sensors are mainly known for their use in video surveillance systems, manufactured by a number of companies (e.g in [9, 8]), to detect motion and provide a trigger to cameras. For their insensitivity to light conditions, in [17] PIR sensors are used to provide a trigger event in a motion-detection application mainly based on cameras for tracking events at night, together with a floodlight. The appearance of an infrared radiating body set off the PIR sensor, which turns on the floodlight enabling the cameras to capture clearly an event such as animals passing by an outdoor detected area. Being low-cost, low-power and small form-factor devices, PIR sensors are suitable for wireless sensor networks, where energy consumption, unobtrusiveness and easy placement in the environment are critical requirements. In [6], a wireless PIR sensor network is used to detect objects and humans for security applications and provide an estimation of the direction of movements. The network is implemented using Mica2 [10] nodes and data gathered by a base station. Tracking algorithms are implemented on the nodes and speed calculation provided accurately, even if influenced by the orientation of the sensors. Sensor networks implemented with PIR are useful where privacy must be preserved together with security. In [19] cameras and PIR sensors are deployed respectively in public and private areas, and their information combined to correlate event such as tracking human motion and undesired access or presence in private areas, such as theft. This work demonstrate benefits of reducing camera deployment in favor of PIR sensors and reports results from a survey on 60 people, stating that people consider motion sensors less invasive for their privacy than cameras.

PIR sensors are often combined with vision systems and other kind of sensors in research focused on robot navigation and localization. In [20] a sound source localizer and a motion detector system are implemented on a human service robot called ISAC, with the purpose of redirect the attention of ISAC. The motion detector system use an infrared sensor array of five PIR sensors and it is integrated with the vision system of ISAC to perform real-time human tracking, in a most inexpensive way. Similar use of PIR sensor can be found in [15], where data from multiple PIR sensors, two color cameras and a pair of microwave sensors are collected, processed and fused to track human presence in the vicinity of a robot. A main motivation to use PIR sensors in pair with cameras is their insensitivity to lighting conditions, to avoid robot collision with humans for their safety in many different conditions.

3. PIR SENSOR NETWORK

3.1 Sensor node architecture

As previously discussed in [4], a sensor node is much more complex than a simple sensor. A microcontroller, a transceiver and a power supply (mainly a battery) together with a variety of sensors form an entity capable of collecting events and information from the surrounding area, analyze them and send message or data to other nodes or to user (see Figure 1).

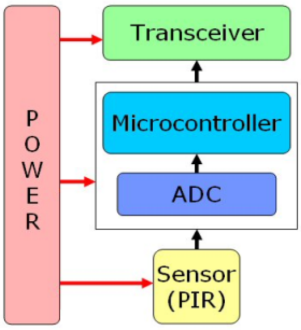

Figure 1: Sensor node architecture

Microprocessor and Transceiver. The hardware used to collect, elaborate and send data from the PIR sensor is the SARD^{TM}Board provided by Freescale^{TM}to develop wireless application over the worldwide free 2.4 GHz ISM band. This development board includes a microcontroller of the HCS08 family (MG9S08GT60), an MC13192 2.4GHz transceiver, a set of I/O pins and a RS232 port to interface the CPU with the external world and several leds and buttons. This development board is very flexible and many types of sensors can be connected to it. Moreover, the microcontroller has the ability to operate in low power mode to save energy and to be waken up by external interrupts generated by the sensor output conditioning circuits. The transceiver is designed to be low-power, short range and fully compatible with the IEEE 802.15.4 standard. It communicates with the microcontroller via a four-wire SPI interface and an interrupt request output. The transceiver supports 250 kbps O-QPSK data in 16 5.0MHz channels and full spread-spectrum encode and decode. It enables communication with other devices up to distances of 40m outdoor.

Sensing element. The sensing element used in our application is a passive infrared (PIR) sensor. Such devices are able to transduce incident infrared radiation into current. Commercial PIRs are sold in pairs with opposite polarization. Such configuration makes the sensor able to detect variations of incident infrared radiation irradiated by bodies moving inside his coverage area which are not at thermal equilibrium with the environment. PIR are used with Fresnel lens to enlarge and shape their area of sensitivity. They are passive sensors with minimal power consumption, ideal for battery powered systems. The PIR sensor used in this work exhibits high sensitivity and reliable performance, high stability to thermal change and high immunity to external noise. By suitably shading its Fresnel lenses we were able to obtain cone of coverage with a vertical angle of 60 degree and an horizontal angle of 38 degree. A single PIR sensor can detect direction of movement. Figure 2 shows the signal detected by sensor when a person passes through the area under control from left to right (the first peak is negative) and from right to left (first peak is positive).

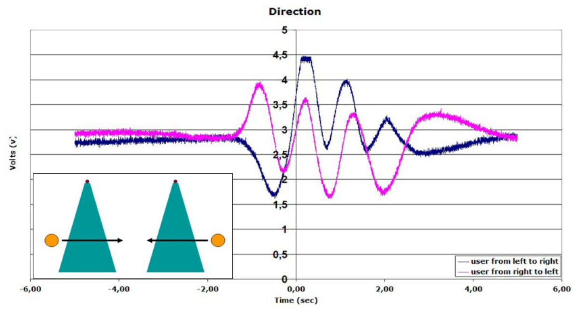

Figure 2: Signals detected by sensor: when a person passes through the area under control from left to right the first peak is negative and from right to left the first peak is positive

In our setup, the typical node include two PIR sensors to enhance the information captured as explained in the following paragraph.

Power Supply. The system can be powered directly by main power or alternatively with a commercial battery at 9V. Internally, the voltage is stabilized at 3.3V.

3.2 Sensing and Acquisition Software

Figure 3 shows the processing data flow from event acquisition to generation of the packet which will be sent by the wireless node. This section describes the role of the different parts.

We are interested in detecting precisely presence and direction of movement, but also more complex movements such as changes in direction within the covered area. In fact, these are information that can be exploited by the vision system for enhancing the accuracy of the video surveillance application, in which presence and direction of movement (of people) are key information.

As outlined above, we augment the information produced by a single node by using 2 PIR sensors (Figure 4(a)) per node. The typical sensors output when a person is walking through the sensor area is the one presented in Figure 4(b). The signal collected by the sensors is digitally converted to be processed by the microcontroller. When a person crosses the monitored area each of the two sensors generates a waveform similar to the one in Figure 2 depending on the direction of movement. We consider interesting events those stimulating a significant variation of the signal (Figure 4(b)): when the input coming from the digital converter exceeds a lower or an upper threshold, a trigger is generated to start the processing algorithm in charge of extracting information from the signal. The analysis, as mentioned above, is aimed at understanding the direction of a person walking in the covered area. Assuming that one person is moving

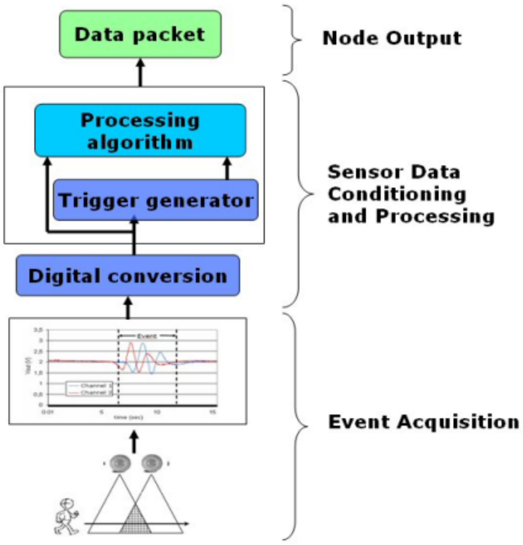

Figure 3: Event acquisition and Sensor Data Conditioning and Processing

from left to right as in Figure 4(a) he will be detected first by PIR_i then by both PIR_i and PIR_j and at last only by PIR_j as it is lightened in Figure 5. In general, a different activation sequence can help identifying changes in direction of movement within the area covered by the array of sensors. Results from the processing is a message containing information about the presence and/or direction of movements in the selected area. The format of the packet is described in the following section.

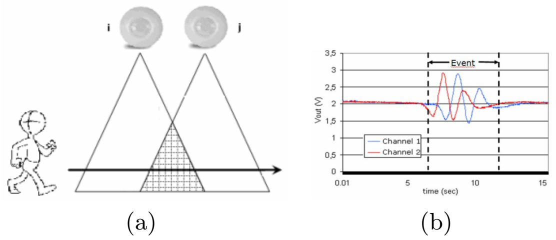

(a) (b)

Figure 4: Sensor node composed by two PIRs

Note that the trigger generator is disabled for a period to be set depending on the application after the detection of an event, avoiding redundant information to be sent. In our case the period is set at 2 seconds. This choice has been verified as not influencing correct analysis, because it does not cause loss of events.

3.3 Network Architecture

Communication Protocol. Communication among nodes is based on IEEE 802.15.4 protocol for wireless networks. This protocol has been designed for applications with limited power budget and relaxed throughput requirements. Its main objectives are ease of installation, reliable data transfer, short range operation, extremely low cost and a reasonable battery life, while maintaining a simple and flexible protocol. This protocol defines two types of devices: full function device (FFD) and reduced function device (RFD). It also defines three possible roles for the nodes of the network: coordinator (which is unique for a network), router and end device. A coordinator and a router must be FFD while an end device may be either FFD or RFD. Usually routers and coordinators are main powered while end devices are located on the field and are battery powered. Two or more devices constitute a wireless personal area network (WPAN). However one device must be a coordinator.

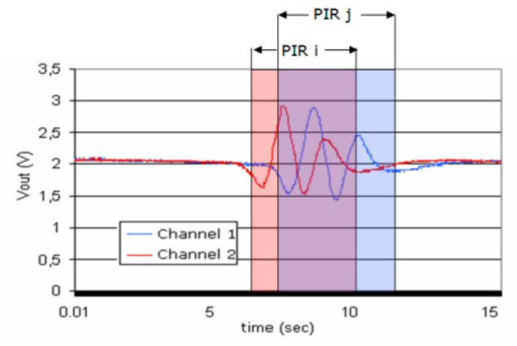

Figure 5: Activation sequence.

Network topology and organization. The network has a star topology, i.e., all the nodes are end devices and communicate only with a central one, the coordinator. The central node (bridge) collects data from the sensor nodes and sends them to another node (sink) which communicate through its RS232 interface to a PC (see Figure 6). Hence, in our application the bridge is the network coordinator while the other nodes are end devices. The sensor nodes, which are located in the courtyard, are battery-powered while the bridge and the sink are main-powered. This topology is suitable to the characteristics of the monitored area. In fact, the sensor are located in a courtyard outside the building while the PC, due to privacy issues, is locked inside a small room, which must be kept closed within the building. Some tests shown that only the sensors close to the door of the building are able to communicate with a device inside the room, while all the courtyard can be covered by a receiver located close to the door.

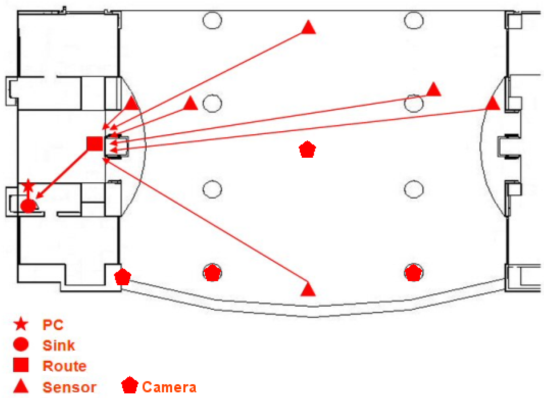

Figure 6: Map of our test bed system

Message format and set of commands. As already men-

Information	Code	Values	Code
Presence	1	Present	1
		Area free	16
Direction	2	From PIR_i to PIR_j	48
		From PIR_j to PIR_i	192

Table 1: Examples of adopted codes

tioned, the information collected by the sensors are sent to the video processing server via RS232 cable. We decided to use an asynchronous communication, that is, the sensor network send data to the server as soon as it collects them. The structure of the messages is shown in Figure 7.

Figure 7: Communication protocol between nodes and sensor manager

Each message is made up of a start byte (the ASCII code 'I'), a sensor ID, an area ID, an indication of length (the number of following couples name-value), several couples name-value and a stop bit (the ASCII code 'F'). Start and stop bit are used for synchronization. Area and sensor node ID are used to uniquely identify the node. The couple name-value encodes the information provided by the sensor. Some examples are reported in Table 1.

4. VISION SYSTEM

4.1 Single Camera processing

Many approaches to people detection and tracking by single cameras have been proposed in the literature. Their schemes are often similar: first, perform motion detection by separating points belonging to still parts from points belonging to moving parts (by means of background suppression, frame difference, or statistical analysis); then, blob analysis aims at grouping spatially correlated points into objects and characterizing them by visual features and motion components; eventually, moving objects are tracked with the aim of keeping track of their identity to further analyze the behaviour.

Our approach from single camera follows this scheme, and it is composed by two main modules: segmentation and tracking. The first module aims at extracting *visual objects*, that are entities that we are interesting in and that we want to analyze separately with respect to the background. Normally, the *visual objects* are objects detected in motion for almost one frame. To this aim background suppression techniques are often adopted and operate by subtracting the current background model B^t from the current frame I^t. The points are extracted and grouped with a labelling process into a set FO^t of foreground objects at instant time t. This set contains both relevant objects and other outliers, such as shadows and noise. To identify shadow points we used a deterministic approach, proposed in [2], based on the assumption that shadows have similar chromaticity but lower brightness than the background on which they are cast.

Objects in the set FO^t considered too small are discarded as noise. The set VO^t of visual objects obtained after the size-based validation is processed by the tracking module that computes for each frame t a set of tracks $T^t = \{T_1^t, ..., T_m^t\}$.

In the case of people tracking, the basic tracking approaches (based on directional rules, or Kalman filters) are not suitable, since humans undergo to deformation in the shape, move with unpredictability and sudden changes in the main direction, and are likely to be occluded by objects or other people. For these reasons, we proposed a probabilistic and appearance-based tracking algorithm able to manage also large and long-lasting occlusions [3]. Despite its accuracy, our tracking fails in the case the person changes his direction when occluded, since the algorithm relies on the hypothesis of constancy of motion during occlusions (being any other hypotheses not reasonable). Since in absence of visibility cameras are useless, this is a concrete and interesting example in which PIR sensors can be useful, as it will be shown in the following section.

The knowledge about VOs and their status is exploited by a selective background model [2] in order to be both reactive to background changes and robust to noise. Selective update is obtained by, on the one hand, not considering moving pixels in the updating process, and, on the other hand, forcing inclusion of stopped objects (previously moving) into the background model. Unfortunately, the system sometimes misclassifies moving objects (such as a person) with stopped objects (such as a door that has been opened). In these cases, the lack of enough resolution prevents the vision system to work properly and PIR sensors might help.

Eventually, scene understanding is a high-level module and heavily depends on the specific application. In the case of video surveillance of people, it includes a posture classification module [3], capable to discriminate between four postures (standing, sitting, crouching, and laying) and, consequently, to detect interesting events, such as a person's fall.

4.2 Consistent labeling

Real video surveillance setups often require multiple cameras, both to cover wider areas and to solve occlusions by exploiting multiple viewpoints. The goal of the consistent labeling and multicamera tracking module is the detection of correspondences between people extracted from each single camera tracking module, and then the computation of a list composed by the best views (selected from the different cameras) of people present in the scene. This list is the input of higher level tasks, as posture classification, face detection, and recognition. We propose an approach of consistent labeling based on geometrical features and homography transformations. For two overlapped cameras C_i and C_j, we compute with a training procedure the homography transformations that bind the ground planes. Full details can be found in [1]. Differently from other methods that check consistency only when objects pass through the edges of the field of views (camera handoff), we compute the assignment each time a new object is detected in the camera C_i in the overlapping. In this case its support point is projected in C_j by means of the homographic transformation. The coordinates of the projected point could not correspond to the support point of an actual object. Thus, we select for the match the object in C_j whose support point is at the minimum Euclidean distance in the 2D plane from these coordinates. This approach is an efficient trade-off between classical techniques that verify correspondences at the camera handoff instant only (as in [12]), and complex methods of 3D reconstruction that find correspondences at each frame preventing any real time implementation (as in [16]). Figure 8 gives a bird-view description of the area acquired by three different cameras; this representation is possible due to the homography transformations between different views. The edges of field of view have been superimposed. The people can be detected by one, two, or even three cameras depending on their position. When a person is in the internal part (where three cameras are overlapped), three different views of the same person are available. In Figure 9 an example of consistent labeling is reported: four people are detected from two different cameras while a person (label 35) is only viewed from a single camera.

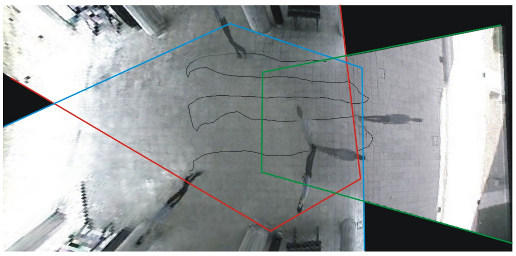

Figure 8: Bird-view description of our test bed environment

Figure 9: Example of Consistent Labeling between two views

5. RESULTS OF MULTI-MODAL INTEGRATION

As stated above, the vision system achieves good accuracy when working alone, but it definitely could benefit from the multi-modal integration from PIR sensors. This section will describe the integration and report some preliminary results. To test the system we have equipped the atrium of our faculty with four cameras and several PIRs, as depicted in Figure 6. Detailed descriptions with particular experimental results of the sensor nodes and of the multicamera system are reported in [4].

The system we implemented is composed by several modules, working in parallel on different threads. In particular, we generate a thread for each camera, devoted to compute the list of people present in the scene exploiting a two stage processing (segmentation and tracking). All the camera threads are tightly connected to a coordinator thread, that detects if the same person is visible in more than one camera, and, in such a situation, it labels the corresponding tracks with the same identifier.

At the same time, a sensor manager coordinates the network of sensors distributed over the monitored area. When a sensor detects an event, a message is sent to the manager. The communication protocol has been reported in Sec. 3.3.

5.1 Sensor-guided Background Update

Figure 10: Opening and closeing doors make unreliable background suppression techniques

Algorithms of motion capture based on background subtraction rely on a very crucial task: the update of the background, specially in presence of illumination changes and moved objects inside the scene. For example, when the doors in Figure 10 are opened, the background scene changes and the detection of people in that area becomes unreliable. To this aim, we use sensors to monitor the area near the doors. If the single camera processing detects a visual object in the door area but the sensors do not capture events, then we assume that the motion is due to an incorrect background. In such a situation we update the background forcing the area covered by the sensor directly with the input image.

In Figure 11 some frames taken from a single camera that is capturing the entrance of our faculty are reported. In the first row are reported the input frames, in the second one the output of the tracking system, and in the last one the background model. Initially (first column) the door was open. Some frames later a person closed the door and from this instant the background becomes inconsistent. In fact, the system erroneously detects the presence of a person in the area of the door (see Figure 11(e)). When the PIR sensor placed near the door does not capture events anymore, the background is updated (last column).

5.2 Detection of Direction Changes during Occlusion

Occlusions are another problem that characterizes video-surveillance systems based on computer vision; for example in the environments of Figure 10 people can walk behind the columns, and, in such a situations, the system is likely to lose them. To face this problem we have introduced some rules inside the tracking system. When a track disappears, it is not deleted immediately, but its appearance is kept unchanged and an estimation of the track position is computed exploiting a constant speed assumption. If the person returns visible again with a similar appearance and a position near to the predicted one, then the system assigns the same label of the disappeared track. However, if the person changes direction during the occlusion, the system is not able to correctly assign the label anymore.

For this reason, we exploit a PIR sensor node placed behind the column. As above mentioned, these sensors detect

(a) Input - frame 1076 (b) Input - frame 2155 (c) Input - frame 2156

(d) Output - frame 1076 (e) Output - frame 2155 (f) Output - frame 2156

(g) Background - frame 1076 (h) Background - frame 2155 (i) Background - frame 2156

Figure 11: Sensor-guided background update.

not only the presence of a person, but also his direction. Then, we can detect a change of direction capturing couples of opposite direction events sent in a short temporal window. In such a situation, we invert the direction of the motion applied to the track to estimate the position frame by frame.

In Figure 12 an example of consistent labeling after an occlusion is reported. The person walks behind a column and, during the occlusion, inverts his direction. The computer vision tracking algorithm is not able to solve the consistent labeling because the person reappears too far with respect to the predicted position (computed with a constant velocity assumption). Using PIR sensors, instead, the change of direction could be detected and the estimated track position can be properly updated. Then, when the person reappears, the tracker assigns the same label (24) used before the occlusion (see Figure 12(b)).

6. CONCLUSIONS

Distributed surveillance is a challenging task for which a robust solution, working on a 7/24 basis, is still missing. This paper is meant to propose an innovative solution that integrates cameras and PIR (Passive InfraRed) sensors. The proposed multi-modal sensor network exploits simple outputs from the PIR sensor nodes (detecting the presence and the direction of movements of people in the scene) to improve the accuracy of the vision subsystem.

Two case studies are reported. In the first, the vision system, based on background suppression, fails due to a door that is opened. Since background is not immediately updated, the door is detected as a moving object (resolution is not sufficient to enable a correct motion detection). In this case, a PIR sensor is used to discriminate between the opened door and a real moving person. In the second case study, a person changes its direction when it is occluded by a column. The vision tracking algorithm relies on the constancy of the speed during occlusions and thus fails. A pair of PIR sensors are, instead, used to detect the change in direction and alerting the vision system. The reported results demonstrate that using the integration between PIR sensors and cameras the accuracy can significantly be increased.

7. REFERENCES

[1] S. Calderara, A. Prati, R. Vezzani, and R. Cucchiara. Consistent labeling for multi-camera object tracking. In *Proc. of Int'l Conference on Image Analysis and Processing*, 2005.

[2] R. Cucchiara, C. Grana, M. Piccardi, and A. Prati.

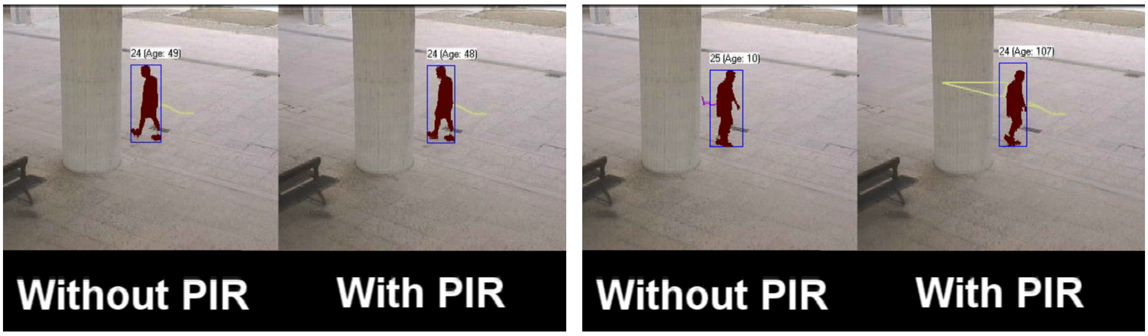

(a) Before occlusion (b) After occlusion

Figure 12: Consistent labeling after an occlusion exploiting a PIR node to detect direction changes

Detecting moving objects, ghosts and shadows in video streams. *IEEE Trans. on PAMI*, 25(10):1337–1342, October 2003.

[3] R. Cucchiara, C. Grana, A. Prati, and R. Vezzani. Probabilistic posture classification for human behaviour analysis. *IEEE Trans. on Systems, Man, and Cybernetics - Part A*, 35(1):42–54, January 2005.

[4] R. Cucchiara, A. Prati, L. Benini, and E. Farella. T_PARK: Ambient intelligence for security in public parks. *Proceedings of IEE International Workshop on Intelligent Environments, Special session on Ambient Intelligence*, June 2005.

[5] D. Culler, D. Estrin, and M. Srivastava. Guest editors' introduction: Overview of sensor networks. *IEEE Computer*, 37(8):41–49, August 2004.

[6] S. de Vlaam. Object tracking in a multi sensor network. *Master Thesis*, August 2004.

[7] G.L. Foresti, C. Micheloni, L. Snidaro, P. Remagnino, and T. Ellis. Active video-based surveillance system. *IEEE Signal Processing Magazine*, pages 25–37, March 2005.

[8] http://www.kalatel.com.

[9] http://www.smarthome.com/7527MC.HTML.

[10] http://www.xbow.com/Products /productsdetails.aspx?sid=3.

[11] C. Jaynes. Multi-view calibration from motion planar trajectory. *Image Vis. Comput.*, 22(7):535–550, July 2004.

[12] S. Khan and M. Shah. Consistent labeling of tracked objects in multiple cameras with overlapping fields of view. *IEEE Trans. on PAMI*, 25(10):1355–1360, October 2003.

[13] J. Krumm, S. Harris, B. Meyers, B. Brumitt, M. Hale, and S. Shafer. Multi-camera multi-person tracking for easyliving. In *Proc. of IEEE Intl Workshop on Visual Surveillance*, pages 3–10, 2000.

[14] J. Li, C.S. Chua, and Y.K. Ho. Color based multiple people tracking. In *Proc. of IEEE Intl Conf. on Control, Automation, Robotics and Vision*, volume 1, pages 309–314, 2002.

[15] Yucong Lu, Lingqi Zeng, and Gary M. Bone. Multisensor system for safer human-robot interaction. *IEEE International Conference on Robotics and Automation*, April 2005.

[16] A. Mittal and L. Davis. Unified multi-camera detection and tracking using region-matching. In *Proc. of IEEE Workshop on Multi-Object Tracking*, pages 3–10, 2001.

[17] H-W. Braun P. Bryant. Some applications of a motion detecting camera in remote environments. *Technical Report*, February 2003.

[18] F.M. Porikli and A. Divakaran. Multi-camera calibration, object tracking and query generation. *Proc. of IEEE Intl Conference on Multimedia and Expo*, 1(1):653–656, July 2003.

[19] S. Rajgarhia, F. Stann, and J. Heidemann. Privacy-sensitive monitoring with a mix of ir sensors and cameras. *Proceedings of the Second International Workshop on Sensor and Actor Network Protocols and Applications*, pages 21–29, August 2004.

[20] A.S. Sekmen, M. Wilkes, and K. Kawamura. An application of passive human-robot interaction: Human-tracking based on attention distraction. *IEEE Transaction on Systems, Man, and Cybernetics*, 32(2):248–259, March 2002.

[21] Z. Yue, S.K. Zhou, and R. Chellappa. Robust two-camera tracking using homography. In *Proc. of IEEE Intl Conf. on Acoustics, Speech, and Signal Processing*, volume 3, pages 1–4, 2004.