A multi-modal approach for activity classification and fall detection

José Carlos Castillo [a] , Davide Carneiro [b] , Juan Serrano-Cuerda [a] , Paulo Novais [b] , Antonio Fernández-Caballero [a] & José Neves [b]

[a] Instituto de Investigación en Informática de Albacete, Universidad de Castilla-La Mancha, Albacete, 02071, Spain

[b] Departamento de Informatica, Universidade do Minho, Campus of Gualtar, Braga, 4710-057, Portugal
Version of record first published: 02 Apr 2013.

PLEASE SCROLL DOWN FOR ARTICLE

Taylor & Francis
Taylor & Francis Group

A multi-modal approach for activity classification and fall detection

José Carlos Castillo[a,*], Davide Carneiro[b], Juan Serrano-Cuerda[a], Paulo Novais[b], Antonio Fernández-Caballero[a] and José Neves[b]

aInstituto de Investigación en Informática de Albacete, Universidad de Castilla-La Mancha, 02071-Albacete, Spain; bDepartamento de Informatica, Universidade do Minho, Campus of Gualtar, 4710-057 Braga, Portugal

The society is changing towards a new paradigm in which an increasing number of old adults live alone. In parallel, the incidence of conditions that affect mobility and independence is also rising as a consequence of a longer life expectancy. In this paper, the specific problem of falls of old adults is addressed by devising a technological solution for monitoring these users. Video cameras, accelerometers and GPS sensors are combined in a multi-modal approach to monitor humans inside and outside the domestic environment. Machine learning techniques are used to detect falls and classify activities from accelerometer data. Video feeds and GPS are used to provide location inside and outside the domestic environment. It results in a monitoring solution that does not imply the confinement of the users to a closed environment.

Keywords: activity classification; fall detection; behavioural analysis

1. Introduction

Every developed country is currently dealing (or about to deal) with the effects of population ageing, which can be defined as a shift in the distribution of a country's population towards greater ages (Weil 2006). It is generally due to lower birth rates and higher life expectancy, resulting in a bigger slice of the elder population. However, population ageing means much more in its consequences (Turner, Giorno, Serres, Vourc'h, and Richardson 1998). There are evidently both social and economical costs associated (Tai and Williams in press). As the situation evolves, economical costs become too expensive to be supported by the diminishing active population, leading to the bankruptcy of the social security systems. In fact, according to a report by the United Nations (2002), not only the working population is declining but also the labour force participation of the older population has declined worldwide over the last decades (Malmberg 1996). According to the same report, the number of older persons has tripled over the last 50 years and will more than triple again over the next 50 years. This phenomenon is happening all over the globe in both developed and developing countries, being some African and South-American countries the only exceptions. However, this challenge for the twenty-first century is also social. Elderly frequently have special needs and require a close and personalised monitoring, mainly due to health-related issues (Dosi 1998). The elder frequently moves to a relative's house or to a care centre (Costa et al. 2009), being out from

their environment, their routines and their life (Al-Hamad, Flowerdew, and Hayes 1997).

People should be allowed to age actively and maintain their quality of life while dealing with all the diseases and limitations that arise (Mandal and Sairam 2012; Polat 2012). This vision, as noted by the World Health Organization (WHO), defines the process of active ageing as optimising opportunities for health, participation and security to enhance quality of life as people age (WHO 2002). Society, as well as the elderly, can benefit from the latter active participation in cultural, civic or spiritual affairs. However, in order to experience this integration, the elderly need to feel independent and safe. In this sense, technological solutions may play an important role. While technological solutions aimed only at supporting the health care professionals in the past, there is a growing trend towards decentralised approaches which provide support to patients in a personalised way, even outside of the hospital environment.

This is being much driven by the concept of Ambient Intelligence (AmI) as defined in 1999 by the ISTAG – Information Society Technologies Advisory Group (ISTAG 1998). AmI describes a potential future in which users are surrounded by 'traditional' devices with embedded computational power communication capabilities that create a sensing and acting network aimed at supporting the user in his/her daily activities or in maintaining a given level of comfort, quality of life and security. Early approaches focused on monitoring patients with cardiac risk (Gouaux

*Corresponding author. Email: josecarlos@dsi.uclm.es

et al. 2002), domestic care and independent living of the elderly (Riva 2003) and learning the user's frequent behaviours in the environment (Aztiria, Izaguirre, Basagoiti, Augusto, and Cook 2010). In an approach in line with this recent vision, our research teams are focusing on the classification of activities and detection of falls in the context of senior care. We chose this specific issue because falls have a significant weight on the worsening of an already precarious health condition. After a fall, a timely detection and rapid response is of utmost importance to minimise its potentially negative effects.

Thus, in this paper a technological approach to detect falls is introduced. We are aware of the challenges that this field of research poses, namely in terms of the quality of the information, the accuracy of the classification and the issues related to privacy. In that sense, a multi-modal approach is followed. A chest-worn accelerometer is used to classify activities; and a video camera is used to detect the location of the user inside the house or a global positioning system (GPS) receiver to provide location outside the home. Some major drawbacks of the use of these approaches can be pointed out:

(1) Threshold-based fall detection approaches often have a significant number of false positives when people perform common daily activities such as sitting down or getting up.
(2) Many accurate approaches use several accelerometers or combine them with other sensors such as gyroscopes. However, people need to remember to place them and have to do it correctly.
(3) Placing video cameras in domestic environments raises many issues, namely in what concerns privacy.

In that sense, we are merging these two approaches to tackle some of the known criticism. Namely, we use a single accelerometer placed in the chest since we can accurately detect falls and classify activities with it (as described in Section 6). However, there is the problem with false positives. In order to address it, we make use of the video cameras, which provide contextual information about the user, namely concerning his/her precise location, trajectory and surroundings (e.g. in which room of the house the user is; where inside the room the user is; what objects are around the user (e.g. sofa, bed)). Based on this, we are able to reduce the number of false positives in the accelerometer and complement it with the context of the user to enable a more accurate interpretation of the data.

Now, focusing on the importance of the user's location, there is no doubt that it is also important to minimise the time spent between the fall and the arrival of help. Hence, fall-detection and activity classification is seen in this work as a location-aware service (Carneiro, Novais, Costa, Gomes, and Neves 2009). As already stated, the lo-

cation of the user is acquired through video feeds inside the domestic environment, which can also provide measures of the level of activity of the user.

While providing location is important for the traditional scenario of users living alone, it is equally important for elderly care centres. In fact, in such institutions it is particularly difficult to monitor every patient without some kind of technological support (Novais, Costa, Carneiro, and Neves 2010). However, the use of video cameras to locate patients has the main disadvantage that the patients must remain inside the environment. In that sense, GPS receivers are also considered so that positions are provided outdoor. Thus, the multi-sensory framework described in this paper provides real-time services that (1) allow the estimation of the location of the users both indoor and outdoor, (2) classify the activities being performed and (3) focus on fall-detection. It may present advantages for elder users living at home as well as for elders living in care centres.

1.1. Related approaches

Extensive work already exists in this field of study, all focused on fall detection and/or activity classification (Riedel, Venkatesh, and Liu 2006; Sivatha Sindhu, Geetha, and Kannan 2012). The first thing these works have in common is that they are all based on relatively new technological solutions. In fact, all these works point out that new technologies can be part of the solution for the problem of fall-detection.

The simplest approach and one of the most common ones consists in using tri-axial accelerometers to measure the magnitude and direction of the acceleration in three orthogonal directions, while the body of the user moves. Generally, an alarm is raised when the value of the acceleration or velocity goes over a given threshold. Different strategies exist under this approach, ranging from the use of one single triaxial accelerometer mounted in the waist of the user (Mathie, Basilakis, and Celler 2001) to the use of accelerometers mounted in the hearing aid housing (Lindemann, Hock, Stuber, Keck, and Becker 2005). From such approaches it is possible to extract information about acceleration and velocity. In order to improve the quality of the data acquired, some researchers focused on placing the hardware in different parts of the body to find the one more suited (Kangas, Konttila, Winblad, and Jamsa 2007). Waist and head are often pointed out as the best places for placing accelerometers. Still in order to get more accurate data, other authors focused on collecting data simultaneously on different parts of the body, as described in Prado, Reina-Tosina, and Roa (2002). This approach uses several accelerometers and one threshold for each accelerometer.

The main critique to all these approaches is that such threshold values can be reached with other daily activities such as sitting down abruptly or jumping, thus resulting in a large number of false positives. Supporting such conclu-

sions, 13 accelerometer-based approaches were analysed by Bagalà et al. (2012). In order to address this drawback, accelerometers can be complemented with other techniques or devices. In Luštrek et al. (2011), location sensors and accelerometers are mixed to detect falls. The authors use several radio tags on the clothing of the user that are used to detect the position of the different parts of the body. The authors claim that it improves traditional accelerator-based approaches by up to 40%. The problem is that it requires the user to set up several tags in specific places in order for the solution to be used. Besides this inconvenience, typical problems associated to such approaches are the user forgetting to do so or not being able to accurately do it.

Other approaches mixing several devices (which can be different or of the same kind) include the combined use of gyroscopes, accelerometers, tilt sensors or vibration sensors (Noury et al. 2000, 2003; Qian et al. 2009). However, the same problems are present: users have to be able and remember to place specific hardware on specific places. The alternative is that this hardware is embedded in the clothing. This results more expensive (e.g. each piece of clothing must have the hardware, wearable hardware is more expensive) and is at the moment still not feasible. In order to address such problems, Zhang, Wang, Liu, and Hou (2009) proposed the incorporation of micro sensors on daily using devices, in this case a cell phone. The authors thus incorporated a tri-axial accelerometer in the user's cell phone showing that it can still detect falls accurately and be less intrusive, with the phone being carried on the pocket or around the neck. Although the authors achieved an accuracy of around 80%, this solution applies only if the user carries the phone with him. However, typically, when inside the house a person does not carry a cell phone.

An alternative line of research focuses on the analysis of video feeds in order to detect falls. This approach raises some concerns right away: (1) the privacy of the user is more significantly threatened than with the previous ones and (2) image analysis is a far more complex subject. Different approaches exist to this problem. In some examples location markers are used to reconstruct a 3D representation of the person's position (Nyan, Tay, and Mah 2008). Again, markers must be worn in specific points of the body. This technique often results impracticable for other reasons (e.g. position markers are not visible at all time). This does not result so invasive as no 'real image' from the person is acquired, only a representation. Other techniques do not require markers but need however access to still images of the person, which may be seen as a more severe violation of privacy, particularly in specific spaces of the house (Brulin, Courtial, and Allibert 2009).

Summing up, despite the wide number of existing approaches, some significant problems can still be pointed out, namely: (1) threshold-based values often fail because they consider the values blindly, without context about what the user is doing; (2) many of the projects analysed focus only on detecting falls and not on classifying daily activities that may be important to understand the context of the user and (3) mixed approaches require different pieces of hardware to be placed and worn by the user, who may forget or be unable to do so accurately.

Motivated by the present issues, we propose an alternative approach for monitoring daily activities and detecting falls. It is a mixed one, merging visual analysis with accelerometer data. In the experiments described in this paper, for commodity, the internal accelerometer of the Android platform was used. In the forthcoming version of the solution, it is being replaced by a LilyPad one, based on the ADXL335 accelerometer. The LilyPad is a triaxial accelerometer that can be sewn to clothing and washed. It is placed at the level of the chest and its data are accessed through the use of a smartphone, that collects it and forwards it to a data gateway. While the accelerometer measures the intensity of the movement of the user in different axes, the visual analysis complements this with information about the context and location of the user. This allows for a higher quality of the information acquired as understanding the context and location of the user allows diminishing the number of false positives associated to given activities, such as sitting down or getting up from a sofa or bed.

The two approaches are also complementary in the sense that the use of cameras is not be desirable in certain spaces of the home (e.g. bedroom, bathroom) while the use of the accelerometer is not so problematic. In that sense, when cameras can be used they are used to increase the quality of the information of the system; when they are not, the system relies on the accelerometer alone. A GPS sensor was also included to allow the seamless monitoring of the user, inside and outside his environment. This seeks to increase the sense of autonomy and security of the user: often, users of monitoring solutions start to feel less safe outside the monitored environment, changing their habits to remain more time inside. This is, obviously, counterproductive, as one of the aims is that the user remains active, particularly socially active. This decision thus aims at increasing the sense of security of the users, independently of the location. The envisioned solution was designed to be simple and cost-effective (e.g. video cameras are widely available nowadays, the LilyPad accelerometer costs around 25 USD), but still efficient in classifying activities and falls. The approaches merged, although controversial when considered separately, result interesting in conjunction, tackling each others' drawbacks.

2. The importance of detecting falls

As people get older, their health conditions worsen, bones get weaker and denser (being osteoporosis a frequent condition) and the organism gets more susceptible to infections. For elderly people, falling worsens the health state in several ways, in a short and long-term perspective, affecting

both the physical and physiological spheres of the elderly. The physical short-term effects of a fall include lacerations, hip fractures or head traumas. Studies point out that 20%–30% of people who fall suffer moderate to severe injuries of these types (Alexander, Rivara, and Wolf 1992; Sterling, O'Connor, and Bonadies 2001). These injuries make it harder to live independently and maintain the regular routines, and may even increase the risk of early death. In the year 2000 falls were the most common cause for traumatic brain injuries (Jager, Weiss, Coben, and Pepe 2000). Falls are also pointed out as the main cause of fractures among old adults, commonly affecting the spine, hip, forearm, leg, ankle, pelvis, upper arm and hand (Scott 1990). A fall has also long-term effects on the life of the elder. Psychologically, the elder that falls, even if not injured, is likely to develop a fear of falling again (Vellas, Wayne, Romero, Baumgartner, and Garry 1997). This fear constitutes a psychological limitation to the regular activities, leading to reduced mobility, loss of physical fitness, isolation and solitude. This will, in turn, increase the likeliness of future falls.

Thus, falls represent a serious social problem. However, they encompass also economical consequences. The Center for Disease Control and Prevention maintains a database for assessing the causes and consequences of falls, under the web-based injury statistics query and reporting system[1]. Reports from 2009 point out that 2.2 million non-fatal fall injuries among older adults were treated in emergency departments in the United States only. Of these patients, more than 581,000 were hospitalised. This represents an evident cost for the public health care systems and insurance companies. The same reporting system points out that in the United States, direct medical costs of falls in year 2000 totalled a little over $19 billion (including fatal and non-fatal injuries). Falls are nefarious not only for their consequences but also for the number of old adults affected. A study by WHO reveals that 30% of the people aged 65 years fall at least one time per year, with this number increasing to 50% after the age of 80 (Giannakouris 2008; Suelves, Martínez, and Medina 2010).

Given all these factors, the negative effects of falls on the society and the economy are evident. To minimise these effects one could focus on preventing falls. In order to do so, older adults should be encouraged to exercise regularly (focusing on improving balance and strength), review their medicines (some have side effects such as dizziness), have their eyes checked and make their homes safer by building ramps, grab bars or improving lighting, for example. However, falls will continue to happen. In that sense, we believe that it is paramount to focus on the decrease of the time spent between the fall and the arrival of help. While the elder is laying on the ground, the risks of infections, myocardial infarction, pneumonia or pulmonary thromboembolism increase. Moreover, this period of time has also an effect on the degree of eventual mobility impairments and on the
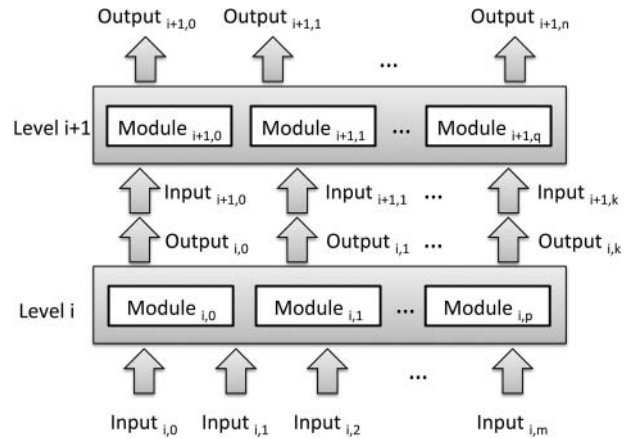

Figure 1. Example of level hierarchy with input/output relations.

recovery time, which in turn aggravates the psychological consequences.

This problem is worsened by the fact that the number of elderly living alone is increasing sharply. According to the UK Office for National Statistics, the number of people aged over 85 rose by 84% between 1981 and 2004, and so has risen the number of this people living alone. Besides, most of these old adults have only the television for company: in a survey of 1018 adults aged over 65 more than half of them did not see their family more than once a month. Such reality calls for automated and autonomous technology-based mechanisms for fall detection that monitor the patient 24/7, independently of the location. Such solutions have the potential to be accurate, to detect a fall in real-time and to minimise significantly the time spent between the fall and the arrival of help. This minimises the negative consequences of the fall and, ultimately, could be the difference between life and death.

3. A framework for human fall detection

In first place, a framework is developed to carry out user monitoring and fall detection tasks (Castillo, Fernández-Caballero, Serrano-Cuerda, and Sokolova 2012; Sokolova, Castillo, and Fernández-Caballero 2012). Several visual sensors implement user detection to assess the position where an important activity is detected. This is particularly important in the event of a fall. To perform fall detection, the framework establishes a set of operation levels with clearly defined input/output interfaces to hierarch the processing. The levels are composed by a set of modules that hold the specific algorithms devoted to each level's operation (see Figure 1). At the lower level, several sensors providing information are found. Following this scheme, one module is included to deal with each sensor. At each level, the framework establishes a set of inputs and outputs available for each level's modules. The inputs and outputs are independent from each other and from the modules. This implies that a module does not need to implement all inputs and
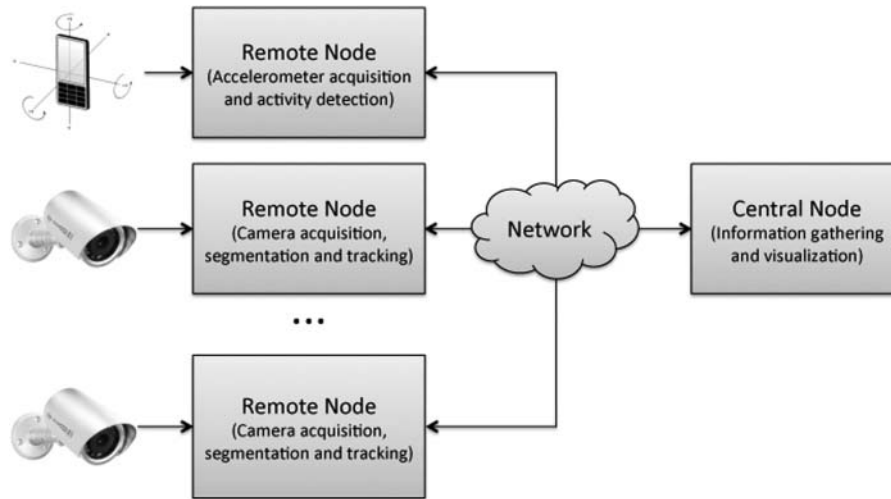
Figure 2. Framework execution scheme. Remote nodes perform information acquisition and processing. The central node gathers and visualises the remote nodes' information.

outputs, but must select its interfaces among them. This is the key for the connection of different kinds of sensors and different algorithms to process their information running in parallel. Following the scheme proposed in Figure 1, the framework levels establish a hierarchy from the sensor acquisition level to the activity analysis one, connecting low level outputs to their respective higher level inputs. The level hierarchy is detailed in Section 3.2.

As aforesaid, the proposed framework is designed to operate with different sources of information coming from cameras and accelerometers. In this proposal, the acceleration sensor is embedded in a mobile phone and connected to the framework through a client/server scheme, acting like another sensor; that is, performing continuous monitoring. Nevertheless, other proposals such as wireless sensor networks might be taken into account to fulfil this purpose.

3.1. Execution model

It is a requirement for a framework in charge of fall detection tasks at home to monitor at different locations, obtaining large amounts of data to be processed. For this reason, the framework execution is defined as hierarchical, providing a series of remote nodes in charge of collecting sensor information and processing it, while a central node is in charge of joining remote information and performing the visualisation. The proposed scheme distinguishes two kinds of remote nodes, namely the one in charge of accelerometer processing and the one in charge of camera processing. The second kind of node may have several instances, each one devoted to one camera. Figure 2 depicts the connection between the framework nodes as well as the execution levels they carry out.

3.2. Levels

In order to optimise the integration of the different algorithms that comprise the framework operation, a stack of four processing levels is proposed here. The lower level corresponds to the acquisition of information coming from the sensors. In this case two sensor technologies are proposed: visual sensors and accelerometers. For this reason, two modules are placed in this first level, dealing with both sensing technologies. Furthermore, this level is in charge of information preprocessing to enhance higher levels' performance. The second level, designated as segmentation, performs the extraction of relevant information from sensor data. A segmentation algorithm isolates the spots corresponding to the objects of interest contained in the images. The spots are also filtered according to a series of constraints to ensure they contain humans. Next, the tracking level adds a temporal component to the segmented spots, matching spots along different captures. This way, the module provides information about the trajectory the objects have followed in the scene. Moreover, tracking algorithms also infer future positions of the objects according to their trajectories. The last level, designated as activity detection, extracts semantic information from the accelerometer information. The data describing the acceleration measured on the handheld device are processed at this level. Specifically, for each instance of the acceleration measured in a given instant t, a seven-tuple is created:

$$acc_t = (x, y, z, m, u, d, t), \qquad (1)$$

where x, y and z denote the values of the acceleration measured on each axis, m contains the module of the acceleration vector, u identifies the user, d the model of the handheld

device and t the time instant in which the acceleration is measured.

The detected activities only involve human as scene objects and are constrained to the classification of falls and common daily movement patterns such as walking, standing or running. Activity analysis from visual information follows a fuzzy-based approach, while activities from information regarding acceleration rely on an implementation of a C4.5 classifier.

4. Human detection from video

The detection of humans using a video feed entails two phases, namely, segmentation and tracking. The segmentation phase consists of a motion-based approach to detect humans in the scene according to their dynamic features. That is, if a human stays still, it is possible that the algorithm produces a miss detection. This problem is solved by the second phase. Tracking uses the information of the detected humans in previous time instants to assess their trajectories, being able to keep them in mind if the segmentation fails. This is also useful in case of actual miss detections, and in this case, the tracking algorithm predicts the future positions of the humans. These two phases are explained in detail below.

4.1. Segmentation

The proposal for human detection in images uses the accumulative computation approach (Fernández-Caballero, López, Castillo, and Maldonado-Bascón 2009). Motion calculation allows obtaining gradually all moving human shapes through this mechanism. The different stages of the algorithm cover the main features of the proposed approach, as depicted in Figure 3 and detailed as follows.

Grey level segmentation. The module segments the original image (named as $I(x, y, t)$) into a preset group of N bands. Each of these bands corresponds to a different range of grey levels from 0 to 255. For example, if we use eight grey level bands, the first band covers a range of grey levels from 0 to 31, the second band covers a range from 32 to 63, and so on. At the end of this stage there are N binary images. The pixels whose grey level value in the original image is in the grey level range of a given band will have maximum value on that particular band while the rest of the pixels have a minimum value. These images are named as GLS_i, where $i \in [0..N-1]$.

Permanence. Now, a charge or discharge value calculation due to motion detection is performed. The module has been designed to obtain the accumulated charge $PM_i(x, y, t)$ on a quantisation basis that will memorise the value of the accumulative computation present at time scale t for each pixel (x, y).

At each pixel (x, y) we are in front of three possibilities: (1) The charge value at pixel (x, y) is discharged down
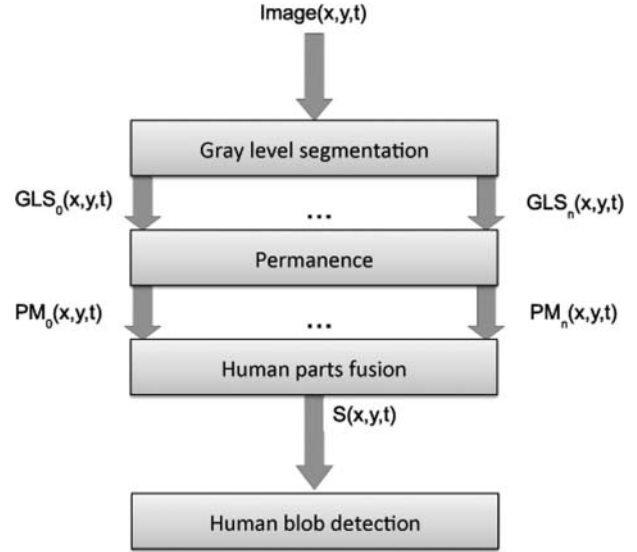

Figure 3.   Accumulative computation approach.

to v_{dis} (the minimum allowed charge value, usually set as 0) when no motion information is detected at band i. No motion information is available as pixel (x, y) does not correspond to band i. (2) The charge value at pixel (x, y) is saturated to v_{sat} (the maximum charge value, usually set as 255 since is the maximum value for a pixel in a 256 grey level image) when motion is detected at t. Motion is detected as the image pixel now belongs to this band at time instant t, and it did not correspond to the band at the previous instant $t - \Delta t$, or vice versa. (3) The charge value at pixel (x, y) is decremented by a value v_{dm} when the pixel keeps on being detected in consecutive intervals t and $t - \Delta t$, meaning that motion information was not detected in that time interval. Of course, the permanence value cannot get off a minimum value v_{dis}. Notice that the discharge of a pixel by a quantity of v_{dm} (usually set as 64 so the pixel will totally be discharged if motion is not detected after four consecutive frames) is the way to stop maintaining attention to a pixel of the image that did capture our interest in the past.

Human parts fusion. During this step, we take the maximum value of all outputs of the N bands to show the detected blobs associated to a moving human as obtained for each colour component:

$$S(x, y, t) = \max(PM_i(x, y, t)), \forall i \in [0..N-1]. \quad (2)$$

An example of the output generated by this stage is available at the centre of Figure 4.

Human blob detection. This module performs a binarisation with threshold Θ_{obj}. Values over the threshold are set to max(255) and below threshold are set to min(0). Once the image is binarised, some morphologic operations, namely a series of erosion and dilation operations, leading
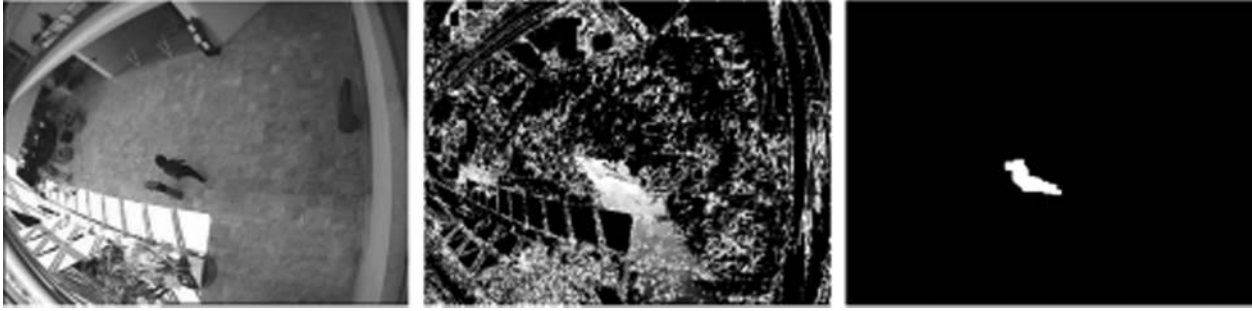
Figure 4. Accumulative computation result. Left, input image; centre, movement detected at different time instants; right, humans detection result.

to eliminate image noise are performed. Finally, spots are filtered based on human characteristic features, such as height, width and compactness. As a result, each human in $I(x, y, t)$ is obtained as a region of interest (*ROI*). A ROI is defined as the minimum rectangle where the human detected can be contained. This rectangle is characterised by its upper-left coordinates $(x_{\min}, y_{\min})$ and its lower-right coordinates $(x_{\max}, y_{\max})$. These ROIs are enlisted into the list of blobs detected L_B.

4.2. Tracking

To deal with humans trajectories, a tracking algorithm was developed (Serrano-Cuerda, López, and Fernández-Caballero 2011). This tracking approach starts from the list of blobs L_B obtained from the segmentation by accumulative computation described above. The tracking algorithm has its own list of ROIs, L_T, which will be refreshed according to the results obtained in each iteration. Firstly, the algorithm will select a ROI from L_B, named as L_{Bi}, $i \in \{0, 1, \ldots, N - 1\}$, comparing it with each ROI present in L_T. The centre of a blob L_{Bi} can be defined as:

$$L_{Bi}.xc = \frac{L_{Bi}.x_{\min} + L_{Bi}.x_{\max}}{2}, \qquad (3)$$

$$L_{Bi}.yc = \frac{L_{Bi}.y_{\min} + L_{Bi}.y_{\max}}{2}, \qquad (4)$$

where $(L_{Bi}.x_{\min}, L_{Bi}.y_{\min})$ will be the initial coordinates from the ROI and $(L_{Bi}.x_{\max}, L_{Bi}.y_{\max})$ its final coordinates as noted before. Thus, those ROIs L_{tj} whose centres are to a distance d lower than an established threshold (which value depends on external factors such as the height where the camera is placed and the dimensions of the room being monitored) will be marked as shown in:

$$d = \sqrt{(L_{Bi}.xc - L_{Tj}.xc)^2 + (L_{Bi}.yc - L_{Tj}.yc)^2}. \quad (5)$$

Now, the ROI with the lowest distance to L_{Bi} will be selected updating its coordinates with the objective of finding correspondence between the human in the ROI of the current frame and the humans contained in previous frames. If a correspondence is not found, a new element is added to L_T with the coordinates of L_{Bi}. In both cases a permanence factor $p(L_{Tj}) = 255$ will be assigned to the new (or the updated) ROI. This factor will be useful to calculate if a human has left the scene. Once the segmented ROIs L_B have been associated with their respective identifiers, a smoothing process must be performed to mitigate the effects from the noise associated to the detection.

A prediction of the possible path will be performed for each ROI that has not been detected in the segmentation, $L'_T = L_T - L_B$. This prediction will be calculated according to the average distance covered between frames, based on $\overline{\Delta xc}$ and $\overline{\Delta yc}$ and its displacement angle. If the permanence value is minimal, the ROI is discarded because it is assumed that it has left the scene. The detected objects may include some noise modifying the shape of the ROI containing them. To mitigate this problem, a smooth operation is applied on the obtained ROIs, based on an adjustment of the height and width with respect to the average values. The detected ROI location will also be modified.

As an example, we will use a possible case where a ROI L_{Tj} has moved in two consecutive time instants. Only the height component $L_{Tj}(t).h$ and its associated formula (see Equation 6) is shown in order to not make the figure too complicated, although the width calculation is performed in a similar way:

$L_{Tj}(t).h$

$$L_{Tj}.h = \frac{N - \lambda}{N} \cdot \overline{L_{Tj}.h} + \frac{\lambda}{N} \cdot L_{Tj}.h, \qquad (6)$$

where N is the number of previous instances of the ROI L_{Tj} used to calculate the average dimensions, λ is the weight of the current height values, and $\overline{L_T.h}$ is the average height.

Afterwards, a location smooth is realised, taking into account the displacement angle θ, as well as the blob centre $(L_{Tj}.xc, L_{Tj}.yc)$ and the newly calculated height and width. The calculations are described in the following equations:

$$L_{Tj}.x\mathrm{max} = \begin{cases} 3/5 \cdot L_{Tj}.w + L_{Tj}.xc, & \text{if } \cos\theta \geq 0 \\ 2/5 \cdot L_{Tj}.w + L_{Tj}.xc, & \text{if } \cos\theta < 0 \end{cases},$$
$$(7)$$

$$L_{Tj}.x\mathrm{min} = \begin{cases} 2/5 \cdot L_{Tj}.w - L_{Tj}.xc, & \text{if } \cos\theta \geq 0 \\ 3/5 \cdot L_{Tj}.w - L_{Tj}.xc, & \text{if } \cos\theta < 0 \end{cases},$$
$$(8)$$

$$L_{Tj}.y\mathrm{max} = \begin{cases} 3/5 \cdot L_{Tj}.h + L_{Tj}.yc, & \text{if } \sin\theta \geq 0 \\ 2/5 \cdot L_{Tj}.h + L_{Tj}.yc, & \text{if } \sin\theta < 0 \end{cases},$$
$$(9)$$

$$L_{Tj}.y\mathrm{min} = \begin{cases} 2/5 \cdot L_{Tj}.h - L_{Tj}.yc, & \text{if } \sin\theta \geq 0 \\ 3/5 \cdot L_{Tj}.h - L_{Tj}.yc, & \text{if } \sin\theta < 0 \end{cases}.$$
$$(10)$$

Now the permanence value will be discharged for each ROI in L_T' whose information was not updated with the segmentation according to the Equation (11) because the belief over the presence of the human in the scene has been reduced:

$$p(L_{T'k}') = p(L_{T'k}') - \delta,$$
$$(11)$$

where δ is a discharge value previously established. Two permanence zones will be defined within a frame, each of them with a different threshold μ_l and μ_h, with $\mu_h \gg \mu_l$. The thresholds for these permanence zones are established according to the situation of the entry points to the scene being monitorised, e.g. mu_l will be set to 0 if the upper limits of the room monitorised belong to a wall. A ROI located near the frame borders will be more easily discarded, because it is very likely for it to have left the scene.

Now the path will be predicted for each ROI in L_T' whose permanence is above the previously mentioned thresholds. This involves the previous calculation of the ROI's covered path, denoted $\overline{\Delta xc}$ and $\overline{\Delta yc}$. The new coordinates are shown in Equations (12)–(15):

$$x_{\mathrm{min}t} = x_{\mathrm{min}t-1} + \cos\theta \cdot \overline{\Delta xc},$$
$$(12)$$

$$x_{\mathrm{max}t} = x_{\mathrm{max}t-1} + \cos\theta \cdot \overline{\Delta xc},$$
$$(13)$$

$$y_{\mathrm{min}t} = y_{\mathrm{min}t-1} + \sin\theta \cdot \overline{\Delta yc},$$
$$(14)$$

$$y_{\mathrm{max}t} = y_{\mathrm{max}t-1} + \sin\theta \cdot \overline{\Delta yc}.$$
$$(15)$$

Finally, the information in L_T will be updated according to the information in L_T'.

5. Activity detection from accelerometer

In this proposal the accelerometer of the Android sensor platform is used to detect falls. The current availability of the platform and the ease of developing applications are key reasons that make use of it. The work started by determining the best position for the device to accurately detect falls. Only healthy subjects participated in the experiments described here. They were requested to perform daily activities such as walking, running or sitting and to simulate falls on a soft ground, in order to avoid injuries. Therefore, the classification algorithms described further ahead were trained using realistic data.

In this process several positions were analysed. When the Android device was in the pocket of the user it did not hold satisfactory results, mainly because the device moves freely inside a loose pocket, generating a significant amount of noise. Afterwards it was attempted to improve the results with the use of a second device placed on the wrist of the user. This neither was effective for detecting falls since wrists move a lot and make it hard to accurately detect falls. After experimenting with different configurations, the decision was in favour of using a single device, placed in the chest of the user. This point has as main advantage to get a relatively low amount of noise related to unimportant movements. However, it effectively senses the acceleration felt in a fall or in other activities such as running or walking. Figure 5 depicts how the Android device is attached to the sweater of the user. The y-axis measures the acceleration due to vertical movements; the x-axis measures the acceleration due to moving sideways and the z-axis measures the horizontal acceleration due to moving forward or backward. In future works a wearable accelerometer that can easily and comfortably be carried by the user will be used.

The differences in the data between the different activities are accessed graphically and through statistical measures of tendency and dispersion. In order to assess the statistical significance of these differences, the Mann–Whitney test is used. This test is a non-parametric statistical hypothesis test for assessing whether one of two samples of independent observations tends to have larger values than the other. The null hypothesis is thus: H_0 = *The medians of the two distributions are equal*. For each two distributions compared, the test returns a *p*-value, with a small *p*-value suggesting that it is unlikely that H_0 is true. We thus compare each axis of each activity with all the others. In all the tests, a value of $\alpha = 0.05$ is used. Thus, for every Mann–Whitney test whose *p*-value $< \alpha$, the difference is considered to be statistically significant, i.e. H_0 is rejected.

Several datasets are recorded with the accelerometer placed as depicted in Figure 5. The datasets contain data about several users performing specific activities for around two minutes: walking, running and standing still. They also include a few falls for each user. Figure 6 depicts an excerpt of a dataset containing data from each of the activities under study. The data for activity *idle* are acquired with the user standing still or performing minor hand movements in an upright position. Data for activity *running* are acquired with the user running forward at a moderate pace. Data
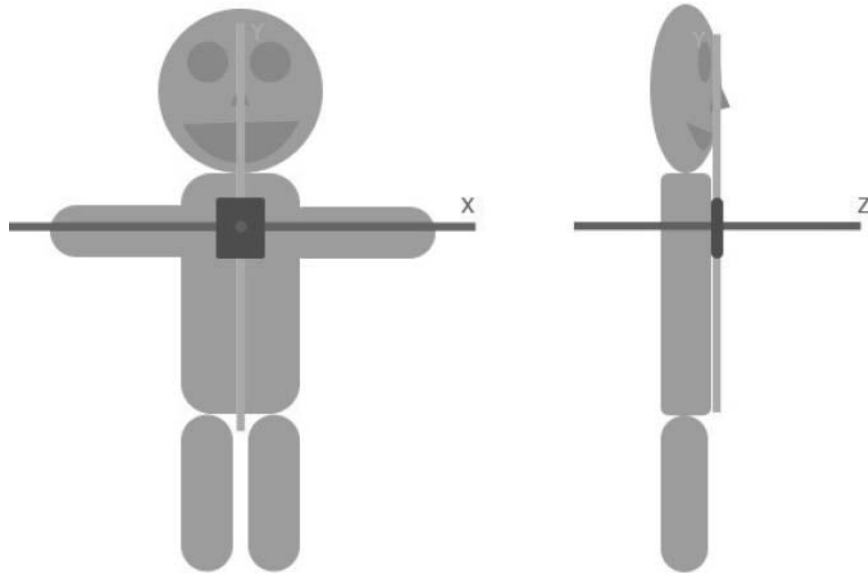
Figure 5. Direction of the three axes of acceleration with the accelerometer placed on the centre of the chest of the user.
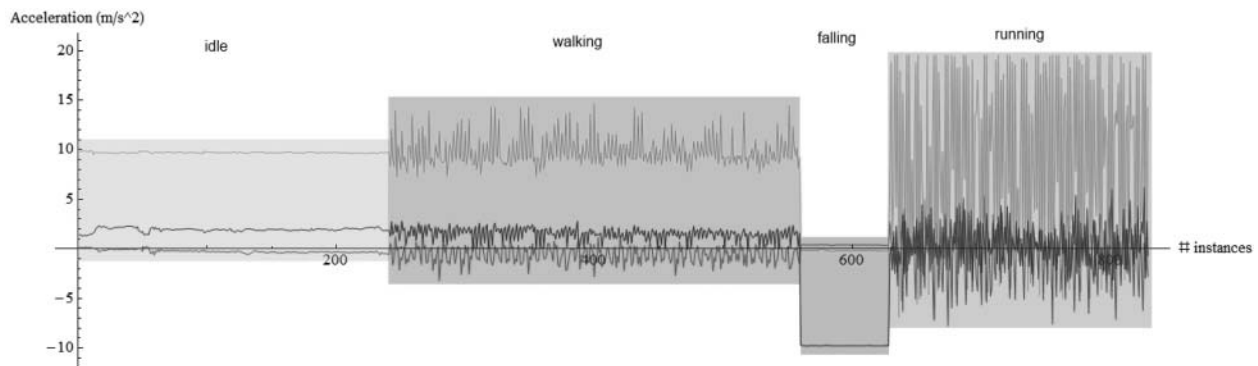

Figure 6. Excerpt from a dataset containing all the activities under study.

for activity *walking* are also acquired with the user moving forward but at a slower pace. Finally, in order to acquire data about falls, each user is requested to fall into a mattress placed in front of him/her. After the fall, the user moves little or remains still.

6. Results

In order to prove the behaviour of the proposed segmentation and tracking algorithms a standard dataset is used (CAVIAR 2004). The dataset sequences provide a source of reference information to compare the performance of the proposed algorithms. The datasets used consist of an indoor environment with a top view camera, able to monitor a hall. To assess the performance of the proposed system in a real environment, *San Juan de Dios* retirement home in Madrid was used as a testbed. It is necessary to point out that falls were performed by actors to ensure elder people safety.

6.1. Metrics definition

Next the metrics used in the experiments are described. Accuracy is the degree of conformity of a measured or calculated quantity to its actual (true) value. The accuracy of an experiment/object/value is a measure of how closely the experimental results agree with a true or accepted value.

$$\text{accuracy} = \frac{\text{true\_positive} + \text{true\_negative}}{\text{true\_positive} + \text{false\_positive} + \text{true\_negative} + \text{false\_negative}}. \tag{16}$$

Sensitivity (of recall) measures the proportion of actual positives which are correctly identified as such (e.g. the

percentage of sick people who are correctly identified as having the condition).

$$\text{sensitivity} = \frac{\text{true\_positive}}{\text{true\_positive} + \text{false\_negative}}. \quad (17)$$

F-score or *F*-measure is a measure of a test's accuracy. It considers both the precision *p* and the recall *r* of the test to compute the score: *p* is the number of correct results divided by the number of all returned results and *r* is the number of correct results divided by the number of results that should have been returned. The *F*-score can be interpreted as a weighted average of the precision and recall, where an *F*-score reaches its best value at 1 and worst score at 0.

$$F\text{-score}$$
$$= \frac{2*\text{true\_positive}}{2*\text{true\_positive} + \text{false\_negative} + \text{false\_positive}}. \quad (18)$$

6.2. Human detection

As shown in Table 1, the quantitative results in comparison with CAVIAR's ground truth are very promising, showing an *F*-score (of *F*-measure) over 91% in all cases, reaching the 97.9% in the best test case scenario. Moreover, the tolerance to false positives and false negatives is also quite high as provided through accuracy and sensitivity measures (97.2 % and 92.7%, respectively).

Table 1. Results of the segmentation algorithm on CAVIAR datasets.

Dataset	Accuracy	Sensitivity	*F*-Score
Browse2	0.885	0.982	0.935
Browse3	0.992	0.855	0.919
Bww	0.995	0.964	0.979
Walk1	0.996	0.915	0.954
Rest1	0.993	0.917	0.953
Mean	0.972	0.927	0.948

Also, some qualitative results are offered to assess the tracking performance. Figure 7 shows the detection and tracking of the moving humans along a sequence from the real scenario. It is depicted how the two humans are perfectly followed through their most accurate positions. A third human (labeled as 3) is only detected at the beginning of the sequence and besides some of the humans are not detected through the different frames. This is caused by the nature of the segmentation algorithm, a movement-based approach, which discards those image zones with little or no movement. The tracking algorithm is in charge of adding the ability to track humans during a period after they stop moving. This is the case observed with human 3 with very little movement, but the tracking algorithm keeps the detection for a few frames.

Figure 7.   Results of the tracking algorithm on the real environment.

6.3. Activity detection

Given the shape of the datasets and the objective of this work, a decision tree constructor is used to classify the instances. Specifically, the `java` implementation of the C4.5 algorithm (Quinlan 1993), also known as J48, is used. The experiments detailed in this section are implemented using the Weka 3.6.3 workbench (Holmes, Donkin, and Witten 1994). We have a particular interest in using decision trees since a model of a decision tree can then be used to classify user activities in real-time and in a real life application by following the explicit rules defined by the model.

The approach consists in developing a different classifier for each axis. The results of the three classifiers are assessed by analysing some performance measures such as the percentage of correctly classified instances and the Cohen's kappa coefficient, which is a statistical measure of inter-rater agreement that tells how much of the accuracy is due to chance. The results of this study are summarised in Table 2.

It is possible to conclude that all the classifiers behave fairly well, not only in terms of correctly classified instances but also in terms of the Kappa statistic, with the classifier of the y-axis performing particularly well. When we compare

these results with the ones of other accelerometer-based approaches, we conclude that they are similar: in Bagalà et al. (2012), 13 fall-detection algorithms are assessed in terms of their performance with real data; they correctly classify, in average, 83% of the falls [$\pm$ 30.3% standard deviation (SD)]. If we compute the average value of the performance of the three classifiers trained with our data, we get the value of 79.57%, which is in line with the algorithms analysed.

The classification of an instance of data as a given activity thus has three different levels of confidence: ranging from *low* (when each classifier points to a different activity for the same instance) to *high* (when the three classifiers point out to the same activity).

6.4. Fall detection

After having acquired the datasets, measures of central tendency and dispersion are determined to gain a better understanding of the data. The main conclusion drawn from this initial analysis is that the data from activities *falling* and *running* are more similar than the remaining ones. Table 3 depicts the values analysed for all activities. Specifically, we consider the maximum, minimum and mean value of the acceleration in each axis for each activity as well as the standard deviation and the median. In a first analysis, the data for each activity have different distributions.

In Table 4 the results of comparing the distributions of the data from all activities (by axis) are depicted. The Mann–Whitney test indicates that most of the activities are indeed significantly different, expect for a few ones. *Running* and *falling* do not have statistically significant differences in the x-axis. This is explained by the fact that the x-axis is the lateral acceleration and, since in the dataset

Table 2. Summary of the performance of the classifiers for the different axes. The total number of instances is 829.

Axis	Leaves	Tree size	Corr.	Incorr.	Pct. Corr.	κ
x	11	21	615	214	74.1858%	0.6297
y	7	13	751	78	90.5911%	0.8655
z	6	11	613	216	73.9445%	0.628

Table 3. Statistics of the datasets of the different activities for each axis.

	Maximum			Minimum			Mean			SD		
	x	y	z	x	y	z	x	y	z	x	y	z
Fall	0.45	0.00	−9.72	0.34	−0.30	−9.83	0.38	−0.18	−9.79	0.02	0.06	0.02
Run	5.28	19.57	6.20	−7.80	−7.08	−3.28	−0.57	8.79	0.45	2.83	9.08	2.20
Walk	1.68	14.67	2.79	−3.21	7.23	−0.26	−0.59	9.84	1.52	0.88	1.56	0.64
Idle	0.26	9.92	2.33	−0.64	9.49	1.29	−0.25	9.72	1.93	0.17	0.06	0.20

Table 4. *p*-value of the Mann–Whitney test comparing the data of all pairs of activities. Results that are not statistically significant for $\alpha = 0.05$ are emphasised.

Activities	x	y	z
Running, falling	**0.0909698**	$1.03786*10^-9$	$2.25063*10^-35$
Running, walking	0.00264749	**0.699423**	$7.44563*10^-10$
Running, idle	**0.174125**	**0.734344**	$8.15983*10^-14$
Walking, falling	$4.5554*10^-25$	$8.48262*10^-39$	$7.78607*10^-39$
Walking, idle	0.000319418	$1.79978*10^-7$	$5.06943*10^-20$
Idle, falling	$4.68619*10^-37$	$4.69001*10^-38$	$5.47099*10^-37$

the users ran and fell forward, the lateral acceleration is very low. For that same reason, the comparison between the distributions of activities *running* and *idle* for the same axis does not hold statistically significant differences. Finally, for the y-axis (vertical acceleration) and for the pairs of activities '*running,walking*' and '*running,idle*' the results also do not point for a significant difference in the distributions.

The differences between the data can also be seen graphically through the Box and Whisker charts that depict the distribution of each of the variables under study, as depicted

in Figure 8. The distribution of the accelerometer data for each activity is shown by axis. Falls are also detected when the user is outside the domestic environment. However, in this case the position of the user is provided by a GPS sensor. It allows to keep track of the user's movements and to react accordingly and more efficiently in case of need when the user is on the move.

A web application that can be accessed remotely by relatives or care givers has also been developed to provide the information about the current user activity and location in real-time. Its interface is depicted in Figure 9. In this
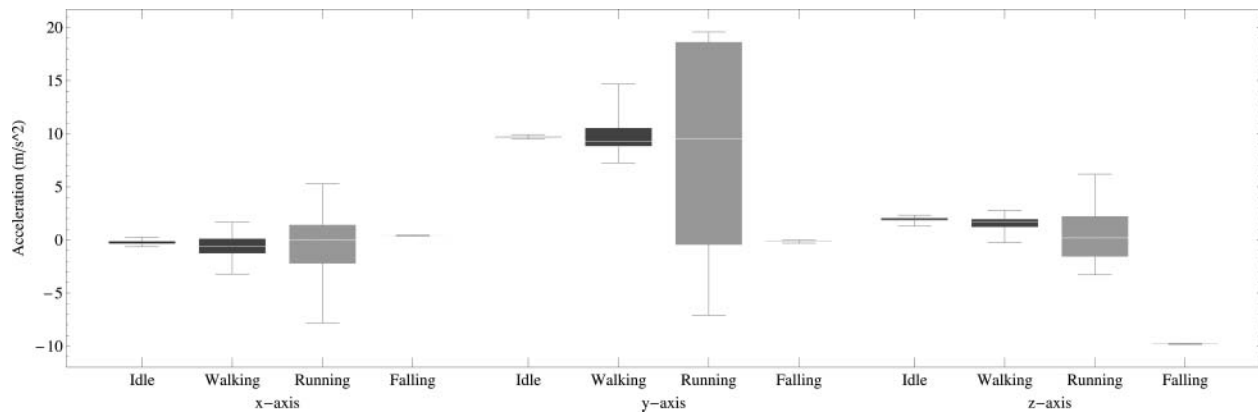

Figure 8. Box and Whisker charts detailing the distributions of the data for each axis and each activity.
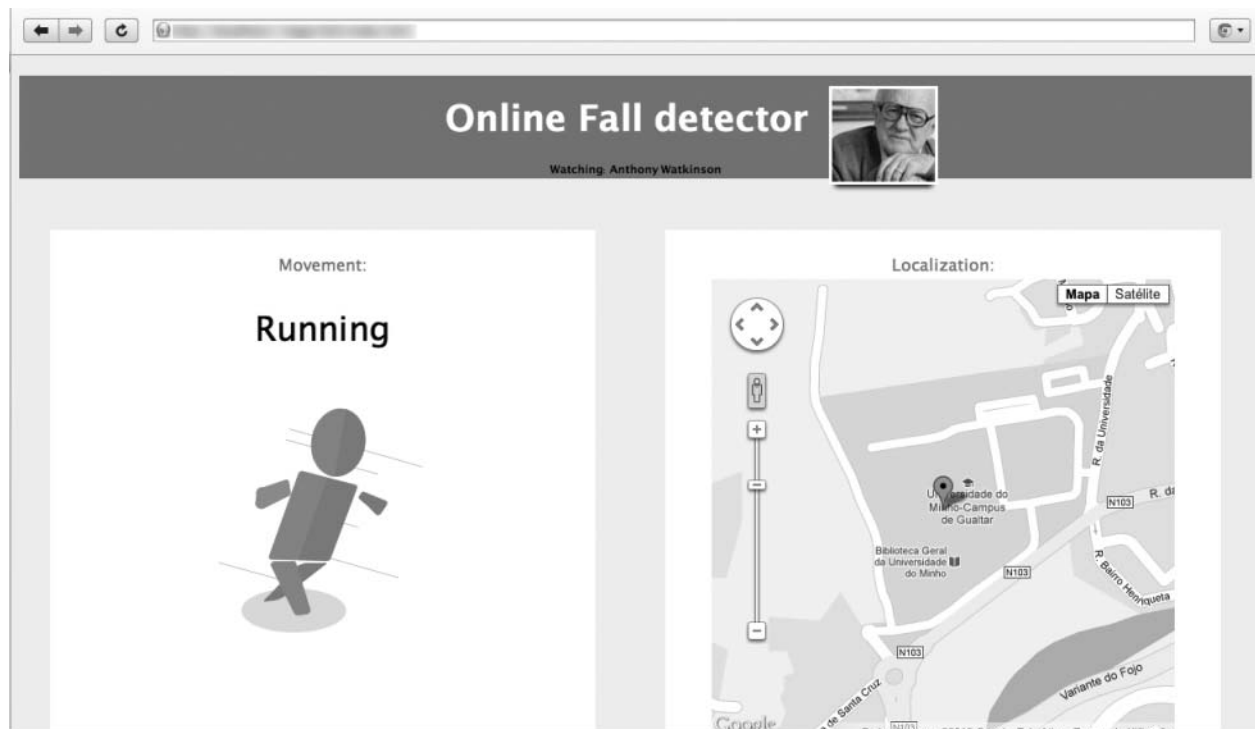

Figure 9. Prototype of the web interface showing the latest known user location and activity.

case the user is on the move, outside the environment being monitored. Therefore, the interface shows the position on the map.

7. Conclusions and future work

The current phenomenon of ageing population poses challenges that call for innovative and sustainable solutions. Specifically, new approaches must be devised to address its social and economical aspects. People should age well by maintaining a certain degree of autonomy and safety, continuing actively involved in their social context. This vision, in line with the concept of Active Ageing defined by the European Union, is frequently threatened by social and economical constraints. In fact, many elderly are alone most of the day, either because they have no informal care givers or because they cannot afford specialised care centres. In our line of research, we believe that technological solutions can in fact address both concerns: provide better care and a sense of security while being affordable.

In this paper the development of a particular solution for activity monitoring and fall detection based on the merging of video cameras, accelerometers and GPS sensors have been introduced. The solution provides the location of the user inside and outside the environment by means of video cameras and GPS sensors. The classification of the activities and the detection of falls is performed through accelerometers. This solution increases the sense of security of elderly living alone in their homes. However, it results even more interesting in the case of specialised elderly care centres with dozens of patients. In fact, it frequently results difficult to efficiently monitor a large number of elder patients. The use of technological approaches as support systems for specialised care givers may improve the safety of the patients. It may improve efficiency of the care provided by increasing the efficiency of detection of events such as falls. This results in a better time management of the staff, since care givers have more free time to dedicate to the patients.

The main innovative aspect of this approach relies in the merger of accelerometers and video cameras. Although these two approaches, when used independently, have some known drawbacks, their joined use may result advantageous. We combine them with a specific purpose: accelerometers classify the activities being performed and video cameras provide the context that allow to correctly interpret such activities and reduce false-positives.

In future iterations we will focus on developing a more user-friendly prototype based on a LilyPad wearable accelerometer placed in the same central position. It will result in a lighter and less invasive solution, making it easier to accept and to be used by the elderly. We will also capitalise on the data concerning the activities being performed. In fact, there is a large interest from specialised care centres in analysing the activities of the patients. On the one hand, we will focus on the elaboration of periodic activity reports detailing the routines of the patients. This way, care givers will use these reports to provide advice for improving the older people lifestyles. On the other hand, we will use machine learning techniques to detect deviations on the normal activity patterns of the patients that indicate a significant change in the health state. This approach will definitely have an impact on the quality of life of the patients, providing an increased sense of security and autonomy, constituting a step forward towards the vision of Active Ageing.

Note

1. The website of the WISQARS™ is available at http://www.cdc.gov/injury/wisqars/ (accessed August 2012).

Notes on contributors

José Carlos Castillo holds an MSc degree in Advanced Computer Technologies (2008) and a PhD degree in Computer Science (2012) from the University of Castilla-La Mancha, Spain. He worked at the natural and artificial Interaction Systems (n&aIS) group at the Albacete Research Institute of Informatics, Spain focusing on computer vision techniques for the detection of human activities and frameworks for intelligent monitoring and activity interpretation. From October 2012, he is working as a postdoctoral researcher at the institute for Systems and Robotics (ISR), Instituto Superior Técnico (IST) of the Tecnical University of Lisbon (UTL), where he is focusing on networked robot systems, robotics and computer vision and intelligent control systems.

Davide Carneiro is a fellow researcher at the Computer Science and Technology Center (CCTC), at the Department of Informatics, University of Minho, Braga, Portugal. He has recently finished his PhD, in which he worked under a joint Doctoral Programme in Computer Science that aggregates three top Portuguese Universities. He has developed scientific research in the field of Artificial Intelligence, namely in Multi-agent Systems, Genetic Programming and Case-based Reasoning applied to the Legal field and Ambient Intelligence. He has participated in several research projects in the fields of Ambient Intelligence and Online Dispute Resolution. He is the co-author of several publications in this field, including book chapters, international journals and conference proceedings. In 2008 he has been awarded the TLeIA08 – a National Award for Artificial Intelligence projects attributed by the Portuguese Artificial Intelligence Association, and in 2009 he has been awarded an Academic Merit Scholarship by the Portuguese Government. His work has been supported

by a research grant and by a PhD Scholarship, attributed by the Portuguese Science and Technology Foundation.

Juan Serrano-Cuerda received his Master in Computer Science from the University of Castilla-La Mancha, Spain in 2007. He holds an MSc degree in Advanced Computer Technologies (2008). His main research interests are computer vision and image processing. He has recently finished his PhD theses on robust human detection based on colour and infrared video fusion in the n&aIS (natural and artificial Interaction Systems) research group at the Albacete Research Institute of Informatics.

Paulo Novais is an Associate Professor of Computer Sciences at the Department of Informatics, in the University of Minho, Braga, Portugal and researcher at the Computer Science and Technology Center (CCTC) in which he is the coordinator of the Intelligent Systems Lab. From the same university, he received a PhD in Computer Sciences in 2003 and his Habilitation in Computer Sciences (Agregação ramo de conhecimento em Informática) in 2011. He is the director of the Master in Informatics Engineering, Master in Informatics and vice director and co-founder of the Master in Law and Informatics at the University of Minho. He develops scientific research in the field of Artificial Intelligence, namely Knowledge Representation and Reasoning, Machine Learning and Multi-Agent Systems, with applications to the areas of Law and Ambient Intelligence. He has leaded and participated in several research projects sponsored by Portuguese and European public and private Institutions and has supervised several PhD and MSc students. He is the co-author of over 100 book chapters, journal papers, conference and workshop papers and books. He is Vice President of APPIA, the Portuguese Association for Artificial Intelligence.

Antonio Fernández-Caballero received his Master in Computer Science from the Technical University of Madrid, Spain, in 1993; and his PhD from the Department of Artificial Intelligence of the National University for Distance Education, Spain, in 2001. He is Full Professor with the Department of Computer Science at the University of Castilla-La Mancha, Spain. He is the Director of the n&aIS (natural and artificial Interaction Systems) research group at the Albacete Research Institute of Informatics. His research interests are in image processing, cognitive vision, neural networks, and intelligent agents. Antonio Fernández-Caballero is Associate Editor of the Pattern Recognition Letters journal. He has authored more than 200 peer reviewed papers.

José Neves is a Full Professor in Computer Science at the University of Minho, in Braga, Portugal, since 1983. Jose Neves is the Deputy Director of the Artificial Intelligence (AI) Group at the same university. He received his PhD in Computer Science from Heriot Watt University, Edinburgh, Scotland, in 1983. His current research interests are in the areas of knowledge representation and reasoning and evolutionary intelligence, aiming to construct dynamic virtual worlds of complex symbolic entities that compete against one another, in order to solve complex problems. Under this setting, fitness is judged by one criterion alone: Intelligence.

References

Alexander, B.H., Rivara, F.P., and Wolf, M.E. (1992), 'The Cost and Frequency of Hospitalization for Fall–Related Injuries in Older Adults', *American Journal of Public Health*, 82, 1020–1023.

Al-Hamad, A., Flowerdew, R., and Hayes, L. (1997), 'Migration of Elderly People to Join Existing Households: Some Evidence from the 1991 Household Sample of Anonymised Records', *Environment and Planning A*, 29, 1243–1255.

Aztiria, A., Izaguirre, A., Basagoiti, R., Augusto, J., and Cook, J. (2010), 'Automatic Modeling of Frequent User Behaviours in Intelligent Environments', in *Sixth International Conference on Intelligent Environments (IE)*, pp. 7–12.

Bagalà, F., Becker, C., Cappello, A., Chiari, L., Aminian, K., Hausdorff, J.M., Zijlstra, W., and Klenk, J. (2012), 'Evaluation of Accelerometer-Based Fall Detection Algorithms on Real-World Falls', *PLoS ONE*, 7, e37062.

Brulin, D., Courtial, E., and Allibert, G. (2009), 'Position Estimation and Fall Detection Using Visual Receding Horizon Estimation', *Proceedings of the 48th IEEE Conference on Decision and Control, held jointly with the 2009 28th Chinese Control Conference (CDC/CCC 2009)*, pp. 7267–7272.

Carneiro, D., Novais, P., Costa, R., Gomes, P., and Neves, J. (2009), 'EMon: Embodied Monitorization', in *Ambient Intelligence European Conference – AmI 2009*, eds. M.Tscheligi, B. Ruyter, P. Markopoulus, R. Wichert, T. Mirlacher, A. Meschtscherjakov, W. Reitberger, LNCS 5859, Springer-Verlag, Berlin, pp 133–142.

Castillo, J.C., Fernández-Caballero, A., Serrano-Cuerda, J., and Sokolova, M.V. (2012), 'Intelligent Monitoring and Activity Interpretation Framework – INT3-Horus General Description', in *16th International Conference on Knowledge-Based and Intelligent Information and Engineering Systems, KES 2012*, Advances in Knowledge-Based and Intelligent Information and Engineering Systems, San Sebastian, Spain, pp. 970–979.

Costa, R., Novais, P., Lima, L., Carneiro, D., Samico, D., Oliveira, J., Machado, J., and Neves, J. (2009), 'VirtualECare: Intelligent Assisted Living', in *Electronic Healthcare*, ed. Dasun Weerasinghe, Springer-Verlag, Series Institute for Computer Sciences, Social Informatics and Telecommunications Engineering, pp. 138–144.

Dosi, G. (1998), 'Sources, Procedures and Microeconomics Effects of Innovation', *Economic Literature*, 26, 1120–1171.

Fernández-Caballero, A., López, M.T., Castillo, J.C., and Maldonado-Bascón, S. (2009), 'Real-Time Accumulative Computation Motion Detectors', *Sensors*, 9, 10044–10065.

Fisher, R.B. (2004), 'The PETS04 Surveillance Ground-Truth Datasets', in *IEEE International Workshop on Performance Evaluation of Tracking and Surveillance (PETS)*, http://homepages.inf.ed.ac.uk/rbf/CAVIAR.

Giannakouris, K. (2008), 'Ageing Characterises the Demographic Perspectives of the European Societies', *Statistics*, 72, 1–12.

Gouaux, F., Simon-Chautemps, L., Fayn, J., Adami, S., Arzi, M., Assanelli, D., Forlini, M.C., Malossi, C., Martinez, A., Placide, J., Ziliani, G.L., and Rubel, P. (2002), 'Ambient Intelligence and Pervasive Systems for the Monitoring of Citizens at Cardiac Risk: New Solutions from the EPI-MEDICS Project', *Computers in Cardiology*, 29, 289–292.

Holmes, G., Donkin, A., and Witten, I.H. (1994), 'Weka: A Machine Learning Workbench', in *Proceedings of the Second Australia and New Zealand Conference on Intelligent Information Systems*, Brisbane, Australia.

Information Society Technologies Advisory Group (1998), *Ambient Intelligence: From Vision to Reality,* For participation in society and business, Information Society Technologies (ISTs) Programme of the European Union Commission.

Jager, T.E., Weiss, H.B., Coben, J.H., and Pepe, P.E. (2000), 'Traumatic Brain Injuries Evaluated in U.S. Emergency Departments, 1992–1994', *Academic Emergency Medicine*, 7, 134–140.

Kangas, M., Konttila, A., Winblad, I., and Jamsa, T. (2007), 'Determination of Simple Thresholds for Accelerometry-Based Parameters for Fall Detection', in *Proceedings of the 29th Annual International Conference of the IEEE EMBS*, Lyon, France, IEEE, pp. 1367–1370.

Lindemann, U., Hock, A., Stuber, M., Keck, W., and Becker, C. (2005), 'Evaluation of a Fall Detector Based on Accelerometers: A Pilot Study', *Medical and Biological Engineering and Computing*, 43, 548–551.

Luštrek, M., Gjoreski, H., Kozina, S., Cvetkovic, B., Mirchevska, V., and Gams, M. (2011), 'Detecting Falls with Location Sensors and Accelerometers', In *Proceedings of the Twenty-Third Innovative Applications of Artificial Intelligence Conference*, pp. 1662–1667.

Malmberg, A. (1996), 'Industrial Geography: Agglomeration and Local Milieu', *Progress in Human Geography*, 20, 392–403.

Mandal, I., and Sairam, N. (2012), 'New Machine-Learning Algorithms for Prediction of Parkinson's Disease', *International Journal of Systems Science*, doi: 10.1080/00207721.2012.724114.

Mathie, M.J., Basilakis, J., and Celler, B.G. (2001), 'A System for Monitoring Posture and Physical Activity Using Accelerometers', In *Proceedings of the 23rd Annual International Conference of the IEEE EMBS*, Istanbul, Turkey, IEEE, pp. 3654–3657.

Noury, N., Barralon, P., Virone, G., Boissy, P., Hamel, M., and Rumeau, P. (2003), 'A Smart Sensor Based on Rules and its Evaluation in Daily Routines', in *Proceedings of the 25th Annual International Conference of the IEEE EMBS*, Cancun, Mexico, IEEE, pp. 3286–3289.

Noury, N., Herve, T., Rialle, V., Virone, G., Mercier, E., Morey, G., Moro, A., and Porcheron, T. (2000), 'Monitoring Behavior in Home Using a Smart Fall Sensor and Position Sensors', In *Proceedings of the 1st International IEEE EMBS Special Topic Conference on Microtechnologies in Medicine and Biology*, Lyon, France, IEEE, pp. 607–610.

Novais, P., Costa, R., Carneiro, D., and Neves, J. (2010), 'Inter-Organization Cooperation for Ambient Assisted Living', *Journal of Ambient Intelligence and Smart Environments*, 2, 179–195.

Nyan, M.N., Tay, F.E.H., and Mah, M.Z.E. (2008), 'Application of Motion Analysis System in Pre-Impact Fall Detection', *Journal of Biomechanics*, 41, 2297–2304.

Polat, K. (2012), 'Classification of Parkinson's Disease Using Feature Weighting Method on the Basis of Fuzzy C-Means Clustering', *International Journal of Systems Science*, 43, 597–609.

Prado, M., Reina-Tosina, J., and Roa, L. (2002), 'Distributed Intelligent Architecture for Falling Detection and Physical Activity Analysis in the Elderly', in *Proceedings of the 2nd Joint EMBS/BMES Conference*, Houston, USA, IEEE, pp. 1910–1911.

Qiang, Li, Stankovic, J.A., Hanson, M.A., Barth, A.T., Lach, J., and Gang, Zhou (2009), 'Accurate, Fast Fall Detection Using Gyroscopes and Accelerometer-Derived Posture Information', in *Sixth International Workshop on Wearable and Implantable Body Sensor Networks*, pp. 138–143.

Quinlan, R. (1993), *C4.5: Programs for Machine Learning*, San Mateo, CA: Morgan Kaufmann.

Riedel, D.E., Venkatesh, S., and Liu, W. (2006), 'A Chemotactic-Based Model for Spatial Activity Recognition', *International Journal of Systems Science*, 3713, 949–959.

Riva, G. (2003), 'Ambient Intelligence in Health Care', *CyberPsychology and Behavior*, 6, 295–300.

Scott, J.C. (1990), 'Osteoporosis and Hip Fractures', *Rheumatic Diseases Clinics of North America*, 16, 717–740.

Serrano-Cuerda, J., López, M.T., and Fernández-Caballero, A. (2011). 'Robust Human Detection and Tracking in Intelligent Environments by Information Fusion of Color and Infrared Video', in *The 7th International Conference on Intelligent Environments, IE'11*, 25–28 July 2011, Nottingham (Reino Unido).

Sivatha Sindhu, S.S., Geetha, S., and Kannan, A. (2012), 'Evolving Optimised Decision Rules for Intrusion Detection Using Particle Swarm Paradigm', *International Journal of Systems Science*, 4312, 2334–2350.

Sokolova, M.V., Castillo, J.C., Fernández-Caballero, A., and Serrano-Cuerda, J. (2012), 'Intelligent Monitoring and Activity Interpretation Framework – INT3-Horus Ontological Model', In *16th International Conference on Knowledge-Based and Intelligent Information and Engineering Systems, KES 2012*, Advances in Knowledge-Based and Intelligent Information and Engineering Systems, San Sebastian, Spain, pp. 980–989.

Sterling, D.A., O'Connor, J.A., and Bonadies, J. (2001), 'Geriatric Falls: Injury Severity is High and Disproportionate to Mechanism', *Journal of Trauma – Injury, Infection and Critical Care*, 50, 116–119.

Suelves, J., Martínez, J., and Medina, A. (2010), 'Lesiones por caídas y factores asociados en personas mayores de Cataluña, España', *Revista Panamericana de Salud Pública*, 27, 37–42.

Tai, G., and Williams, P. (in press), 'Quantitative Analysis of Investment Allocation Over Various Resources of Health Care Systems by Using Views of Product Lines', *International Journal of Systems Science* (in press), doi: 10.1080/00207721.2012.684898.

Turner, D., Giorno, C., Serres, A., Vourc'h, A., and Richardson, P. (1998), 'The Macroeconomic Implications of Ageing in a Global Context', OECD Economics Department Working Papers, No. 193, OECD Publishing.

United Nations (2002), *World Population Ageing: 1950–2050*, New York: United Nations.

Vellas, B.J., Wayne, S.J., Romero, L.J., Baumgartner, R.N., and Garry, P.J. (1997), 'Fear of Falling and Restriction of Mobility in Elderly Fallers', *Age and Ageing*, 26, 189–193.

Weil, D. (2006), *Population Aging*, Working Paper Series, 12147, National Bureau of Economic Research.

World Health Organization (2002), Active Ageing a Policy Framework, Geneva: WHO.

Zhang, T., Wang, J., Liu, P., and Hou, J. (2006), 'Fall Detection by Embedding an Accelerometer in Cellphone and Using KFD Algorithm', *International Journal of Computer Science and Network Security*, 6, 277–284.