

Context-based fall detection and activity recognition using inertial and location sensors

Hristijan Gjoreski^{a,b*}, Matjaž Gams^{a,b} and Mitja Luštrek^a

^a*Department of Intelligent Systems, Jožef Stefan Institute, Jamova 39, 1000 Ljubljana, Slovenia*

^b*Jožef Stefan International Postgraduate School, Jamova 39, 1000 Ljubljana, Slovenia*

Abstract. Accidental falls are some of the most common sources of injury among the elderly. A fall is particularly critical when the elderly person is injured and cannot call for help. This problem is addressed by many fall-detection systems, but they often focus on isolated falls under restricted conditions, not paying enough attention to complex, real-life situations. To achieve robust performance in real life, a combination of body-worn inertial and location sensors for fall detection is studied in this paper. A novel context-based method that exploits the information from the both types of sensors is designed. It considers body accelerations, location and elementary activities to detect a fall. The recognition of the activities is of great importance and also is the most demanding of the three, thus it is treated as a separate task. The evaluation is performed on a real-life scenario, including fast falls, slow falls and fall-like situations that are difficult to distinguish from falls. All possible combinations of six inertial and four location sensors are tested. The results show that: (i) context-based reasoning significantly improves the performance; (ii) a combination of two types of sensors in a single physical sensor enclosure is the best practical solution.

Keywords: Context-based reasoning, fall detection, inertial sensors, location sensors, activity recognition

1. Introduction

Falls are among the most critical health problems for the elderly [6]. Approximately 30% of people over the age of 65 fall each year, and this proportion increases to 40% in those aged more than 70 [30]. About 20% of the elderly who fall require medical attention [10]. Furthermore, falls and the fear of falling are important reasons for nursing-home admission [28]. Falls are particularly critical when the elderly person is injured and cannot call for help. These reasons, combined with the increasing accessibility and miniaturization of sensors and microprocessors, are driving the development of fall-detection (FD) systems.

Even though fall detection has received significant attention in recent years, it still represents a challenging task for two reasons. First, most of the current approaches define a fall as having greater accelera-

tions than normal daily activities. However, since there are several everyday fall-like with high acceleration, such as sitting quickly or lying down quickly, focusing only on a high acceleration can result in many false alarms. Second, not all falls are characterized by a high acceleration. Rubenstein et al. [25] showed that 22% of the falls experienced by the elderly are slow and are caused by dizziness and vertigo (13%), and drop attacks (9%). Therefore, the detection of slow falls should be an intrinsic part of a successful fall-detection system.

To overcome the problems of the existing fall-detection methods discussed above, we propose a new approach to detect falls that combines body-worn inertial and location sensors. The implementation of the approach is named CoFDILS (Context-based Fall Detection using Inertial and Location Sensors) and it uses context information from both types of sensors to determine whether a fall has occurred. It exploits three context components: the user's activity,

*Corresponding author. E-mail: hristijan.gjoreski@ijs.si.

body accelerations and location information. The recognition of the user's activity is the most demanding of the three, so we devoted particular attention to it. We used a machine learning (ML) approach, and the results and the discussion are presented in a separate subsection.

The evaluation was performed on a special real-life scenario that includes fast falls, slow falls and non-fall situations that are difficult to distinguish from falls. In addition, we tested 1023 possible body-placement combinations of six inertial and four location sensors in order to find the best-performing sensor placements for FD and activity recognition (AR) and therefore to achieve the lowest discomfort for the user.

The paper is organized as follows. First, an overview of the related studies on FD is presented in Section 2. In the next two sections, the sensor equipment (Section 3) and the architecture of our system (Section 4) are described. Sections 5 and 6 describe the data preprocessing and AR procedures, respectively. In the next two sections we describe the context components (Section 7) and the context-based reasoning in CoFDILS (Section 8). After that, the experimental setup is presented in Section 9, and the results including the discussion are presented in Section 10. Finally, we conclude the paper and propose directions for future work in Section 11.

2. Related work

FD approaches can be divided into those using non-wearable (i.e., ambient) and wearable (i.e., body-worn) sensors. The most common non-wearable approach is camera-based [14,21]. Although this approach is physically less intrusive to the user compared to the body-worn sensors, it suffers from issues such as low image resolution, target occlusion and time-consuming processing. Probably the biggest issue is the user privacy: the user has to accept the fact that a camera will record him/her.

In recent years, studies that use sound and vibration sensors are gaining attraction. However, these sensors proved to be efficient only when combined with other sensors, especially the wearable inertial sensors [4,5]. We are considering these sensors as future addition to our system, since additional sensing modalities would enhance the context-based reasoning.

Another approach using non-wearable sensors was proposed by Botía et al. [1] and Muñoz et al. [22].

Their system was able to detect most of the alarming situations using three types of sensors: infrared motion sensors, pressure sensors and main door open detector. In the first study, by Botía et al., the authors mainly focused on finding the best time intervals which should be considered in order to raise an alarm. In the second study, by Muñoz et al., the authors proposed an alert management tool for supporting the caregivers in their task of monitoring and validating alerts. The focus of this study is not the accuracy of detection of alarming situation, but the proposed support tool which enables caregivers to easily confirm or dismiss a potential alarming situation. In both cases, they showed that in the case of multiple persons, their system is prone to false alarms, which is an important disadvantage of the systems that use only non-wearable sensors.

Most of the studies on FD and AR are based just on inertial sensors [16]. Usually, they are focused only on fast falls [19,33], which are not difficult to detect using the acceleration signal. The non-fall events used to test for false positives are usually normal, everyday activities [12,21], not events chosen specifically because they are easily mistaken for falls. In contrast, we used complex falls and everyday events that appear like falls. An example where FD was evaluated on events difficult to recognize as falls or non-falls is the work by Li et al. [17]. By applying thresholds to two inertial sensors, they detected a fall with 90.1% accuracy. The recall value of their method on a fall event ending with sitting was 50% and for a non-fall event, quickly lying on a bed, was 40%. By combining one inertial and location sensor, we were able to achieve 99% and 100%, on similar events, respectively.

A combination of inertial and location sensors was described in Zinnen et al. [34]. However, their goal was AR for car-quality control and not FD. Their approach was based on high-level primitives that were derived from a reconstructed human-body model by using inertial sensor data. The location data was mainly used to estimate the person's location near the car. In our approach, beside the location of the user in the apartment, the location features were also used for the recognition of the user's activity.

We are not aware of any prior publication that studies a combination of body-worn inertial and location sensors for FD and AR, except ours [9,20]. The first study, i.e. [9], is a recently published conference paper, which is used as a base for this paper. The most significant changes are the formal definition of the context-based reasoning scheme, and the descrip-

tion of the activity recognition module as a separate task including its results and appropriate discussion. In the second study presented by Luštrek et al. [20], we focused on location-based FD, and we considered only a single accelerometer to detect the impact of the fall and the orientation of the user. The main additions of the study presented here compared to our previous work are: (i) a machine-learning model that recognizes the activity of the user; (ii) a thorough analysis of the system's complexity and invasiveness to the user by analyzing the performance of all the possible body-placement combinations of 10 sensors; and (iii) an explicit presentation of the context-based reasoning algorithm, the core of our system.

A context-based approach to FD is presented in the study by Li et al. [18]. However, they used a different fall-detection method and different types of sensors to extract the context information, compared to our approach. In particular, they used 5 body-worn accelerometers and 2 environmental sensors that monitored the vibration of the furniture. They combined the user's posture information extracted from the accelerometers, and the context information extracted from the environmental sensors, in order to detect the fall situations. Although they also analyzed slow falls and fall-like situations, their evaluation was performed on only 3 test subjects, while we tested our method on 11 subjects. The advantage of our location system, compared to the environmental sensors, is that it provides richer information about the user's situation, e.g., the user's location, the sensor's height, etc. The environmental sensors used in their research can only inform about the presence/absence of the user at a specific location where the sensor is installed. We tested all the combinations of 10 sensors and found a satisfactory performance with single sensor enclosure, while they analyzed only the fixed 5 accelerometer placements on the body.

In general, FD approaches that also exploit the activity of the user tend to be more successful than those relying on high acceleration only. Most of them try to recognize if the user is lying after a potential fall trigger (e.g., high acceleration) [17,24]. Others recognize the fall as one of several elementary activities [4,5,12,19]. There are also some that use the activity information as input to the FD. For instance, to recognize a fall, Sixsmith et al. [26] and Naranjo et al. [23] used two and four levels of activeness, respectively. In our research, the AR is a separate ML component which recognizes not only the level of activeness but the appropriate elementary activity of the user, which is further analyzed by the CoFDILS scheme in order to recognize a fall.

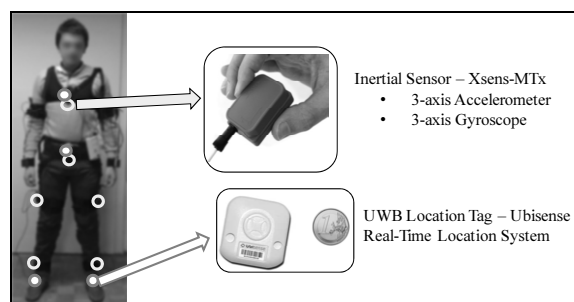


Fig. 1. Sensor equipment. The non-filled circles represent the inertial sensors and the filled circles represent the location tags.

3. Sensor equipment

The CoFDILS sensor equipment consists of inertial and location sensors (Fig. 1). These types of sensors were chosen because inertial sensors are relatively cheap and portable, and location sensors provide rich information about the user without significantly compromising the user's privacy.

Six inertial sensors were placed on the chest, waist, left thigh, right thigh, left ankle and right ankle (non-filled circles in Fig. 1). Since only activities that are associated with the user's legs and torso were studied, the arm- and wrist-sensor placements were not considered. Xsens-MTx inertial sensors [32] were used, but the methods developed in this research are general and can be applied to any type of inertial sensor.

Four location tags were placed on the chest, waist, left and right ankle (filled white circles in Fig. 1). They emit UWB radio signals, which are detected by sensors fixed in the corners of a room, and their coordinates are computed. The location system used in CoFDILS is UbiSense [29]; it is a real-time location system used to track subjects indoors. Note that for simplicity the term sensor is also used for the body-worn location tag.

The data-sampling frequency of the sensors was set to 10 Hz because of the UbiSense's hardware limitations. Although the inertial sensors do not have the same limitation, the data is sampled at the same frequency to simplify the synchronization.

4. System architecture

The architecture of the CoFDILS system is shown in Fig. 2. It consists of three modules: data prepro-

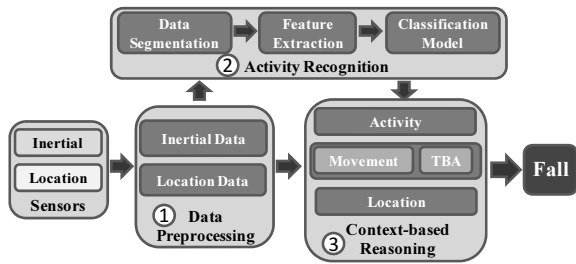


Fig. 2. The CoFDILS system architecture.

cessing (no. 1), activity recognition (no. 2) and context-based reasoning (no. 3).

In the first module, the data from both types of sensors is stored and preprocessed. Next, the flow of data splits into two. On the top, the AR is performed: first, the data is segmented by applying a sliding window, then features are extracted from each data segment, and finally, the constructed feature vector is fed into a classification model, which recognizes the activity of the user. On the bottom, context-based reasoning about the user's situation is performed. This module analyzes the recognized activity of the user and additional context information from the preprocessed data. The motivation is that the context information depends on the type of sensors. Inertial sensors provide body-movement information and the detection of a high-acceleration fall pattern, based on a threshold-based approach (TBA). Location sensors provide the location of the user in the room or in the apartment. The system evaluates the information from various sources in light of its contexts and concludes whether a fall alarm should be issued. Each module in Fig. 2 is presented in more detail in the sections that follow.

5. Data preprocessing

5.1. Inertial data

An inertial sensor provides raw data that consists of accelerations (from an accelerometer) and angular velocities (from a gyroscope) along three perpendicular axes.

The raw data was filtered with low-pass and high-pass filters. The low-pass filter removes the movement of the sensors, which leaves only the gravity component. This information is particularly useful for the assessment of the sensor-inclination angles. In contrast, the high-pass filter removes the gravity and leaves only the sensor movements. These filters were

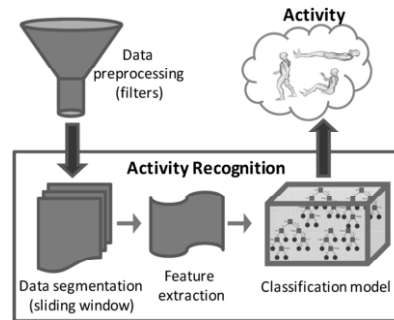


Fig. 3. Activity recognition flow.

applied separately: if the gravity component is needed, the low-pass filtered data is used; otherwise, the high-pass filtered data is used [27].

5.2. Location data

The Ubisense output consists of the 3D coordinates of the sensors that are attached to the user's body. In a typical open environment, the localization accuracy is about 15 cm, but in practice it may occasionally drop to 200 cm or more. Therefore, filtering was performed in order to tackle the problems with the Ubisense system [13].

First, a median filter computed each coordinate as the median of the measured values in a time window. This type of filtering removes large, short-term deviations of a measured coordinate from the true one. Second, the coordinates were corrected with a filter enforcing anatomic constraints based on the user's height and the body proportions. After that, a Kalman filter was used to smooth the data.

6. Activity recognition

To recognize the activities of the user, machine learning (ML) was used [15]. The idea of the ML approach was to learn a classification model that will be able to classify the current activity of the user. The process was divided into three parts: first the continuous sensor data stream was segmented, then for each segment features were extracted, and finally the trained classification model recognized the user's activity.

6.1. Data segmentation

The first step in the ML-based AR is the data segmentation, which uses an overlapping sliding-

window technique, dividing the continuous sensor-stream data into data segments – windows. A window of a fixed size (width) moved across the stream of data, advancing by half its length in each step. Preliminary tests showed that a one-second window size and half-second overlapping was a reasonable trade-off between the duration of the activities and the recognition delay.

6.2. Feature extraction

In the feature extraction procedure, the relevant features were extracted using the preprocessed sensor data in each data window.

6.2.1. Inertial features

This subsection briefly describes the features extracted from the inertial sensors' data and used in the AR [8]. The total number of features per sensor is 25: 8 for the gyroscope data and 17 for the accelerometer data, divided into four groups:

- Statistical features (total 20). The Mean Value and the Standard Deviation were extracted for both the acceleration and gyroscope data; additionally, the Root Mean Square (RMS) was calculated only for the accelerometer data. A feature-selection analysis showed that the RMS was a redundant feature for the gyroscope data.
- Movement intensity feature (AVC feature, explained in Section 7.1.2).
- Sensor inclination angles (total 3). Since most of the time the main component of the acceleration vector was the gravity, they were calculated as the angles between the acceleration vector and each of the axes. For instance, the angle φ_x between the acceleration vector and the x axis was computed as follows:

$$\varphi_x = \arccos\left(\frac{a_x}{\sqrt{a_x^2 + a_y^2 + a_z^2}}\right) \quad (1)$$

where the values a_x , a_y and a_z represent the acceleration vector.

- Difference between the maximum and minimum value of the acceleration vector in the current data window.

6.2.2. Location features

The following features were extracted from the location sensors data:

- The z (height) coordinate of the sensor
- The Euclidian distances between each pair of sensors
- The z -distances between each pair of sensors (difference in heights)
- The Euclidian distances between each pair of sensors in the xy plane
- Two velocity-based features: the first one is the absolute velocity of the sensor, and the second one is computed as the velocity of the sensor in the z direction

6.3. Classification model

Once the feature vector was formed, it was fed into the classification model, which recognized the activity of the user. The classification was performed using the API of the software toolkit WEKA [31]. Among the several methods tested (Decision trees, Naive Bayes, KNN, SVM and Random Forest), Random Forest yielded the best results in preliminary tests [7,11]. Random Forest is an ensemble of decision trees in which the final decision is formed by a majority vote of the tree models [1].

7. Context components

The most important novelty in our fall-detection method (CoFDILS) is the use of the context information. In general, a context is defined as any information that can be used to characterize the circumstances in which an event occurs [3]. In CoFDILS, the context information consisted of three components: (i) the user's body accelerations, (ii) the user's activities and (iii) the location of the user.

7.1. Body accelerations

7.1.1. Threshold-based approach (TBA)

The threshold-based approach (TBA) was used as one of the components in CoFDILS, as well as a baseline for comparison. The rationale for this method was that the acceleration pattern during a typical fast uncontrolled fall (shown in Fig. 4) is a decrease in the acceleration (free fall) followed by a rapid increase (impact with the ground). For our implementation of the TBA, the difference between the maximum and minimum accelerations within a one-second window was calculated. If the difference exceeded the threshold and the maximum appeared

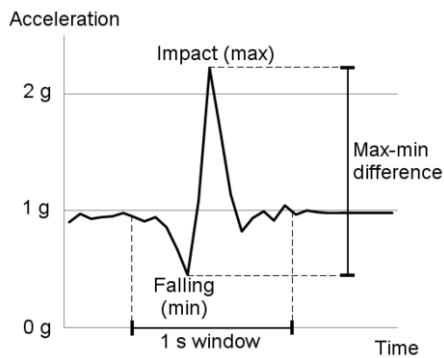


Fig. 4. Acceleration pattern during a fall.

after the minimum, a fall was declared. The threshold was chosen empirically based on preliminary data [7].

7.1.2. Body movement

During motion the accelerometers produce a changing acceleration signal, and the fiercer the motion, the greater the change in the signal. Using these changes a feature was extracted: Acceleration Vector Changes (AVC) [7]. This feature sums up the differences between consecutive values of the lengths of the acceleration vector, and divides the sum by the time interval (one second):

$$AVC = \frac{\sum_{i=1}^n |length_i - length_{i-1}|}{T_n - T_0} \quad (2)$$

T_0 is the time stamp for the first data sample in the window, and T_n is the time stamp of the last data sample. By applying a threshold to the AVC value, the movement of a sensor is detected.

7.2. Activity

Seven elementary activities that can also be interpreted as body postures were studied: *standing*, *sitting*, *lying*, *sitting unusual*, *on all fours*, *going down* and *standing up*. We decided only for these activities because they are the most common, everyday activities and are also the most relevant for the detection of falls and distinguishing them from non-falls.

7.3. Location

The location of the user was provided by the location system, which outputs the 3D coordinates of the location sensors that were attached to the user's body. This way it captured the location of the user in the

apartment and also the height of each sensor. Even though the location sensors provide relatively rich information about the user locations during the day, only the user's presence/absence in locations such as the bed, chair and floor (using also the height of the sensors) was relevant for the FD. Further and more thorough analysis of the location data is not a part of this research.

8. Context-based reasoning

The context-based reasoning schema in CoFDILS is presented in Fig. 5. The general idea is that each of the previously described components uses the information from the other two as context, and reasons about the user's situation. Therefore, there are three possible cases: (i) the *body acceleration* component uses the *activity* and *location* as context; (ii) *activity* uses the *body acceleration* and *location* as context; and (iii) *location* uses the *activity* and *body acceleration* as context.

To explain the basic principle of the context-based reasoning, let us consider the following example in which a user is lying down quickly on a bed, i.e., a non-fall situation. In this case, the acceleration component, i.e., the TBA, recognizes a high acceleration (Case 1 in Fig. 5). If this component reasons by itself, a wrong decision would be formed: a fall would be detected. If the activity of the user is additionally evaluated, the decision would still be wrong (a high acceleration and lying activity = a typical fast fall). However, when the location of the user is evaluated (the bed), the final decision is corrected into non-fall (quickly lying on the bed).

Once we designed the reasoning scheme, we used *context variables* to represent the following three context components: activity, location and body movement. Each context variable contains the value of the component at each moment in time. Table 1 shows the context variables and their possible values. Since the activity data was segmented using one-second window size with half second overlap (Section 6.1), we used the same sampling frequency for the other two context variables. Therefore, each context variable was updated with a new value twice in a second.

The TBA component is a special case because it is instantaneous, happens at a specific moment in time and does not last. Therefore, it was used simply as an independent indicator of an alarming situation which is checked in the context of the other components.

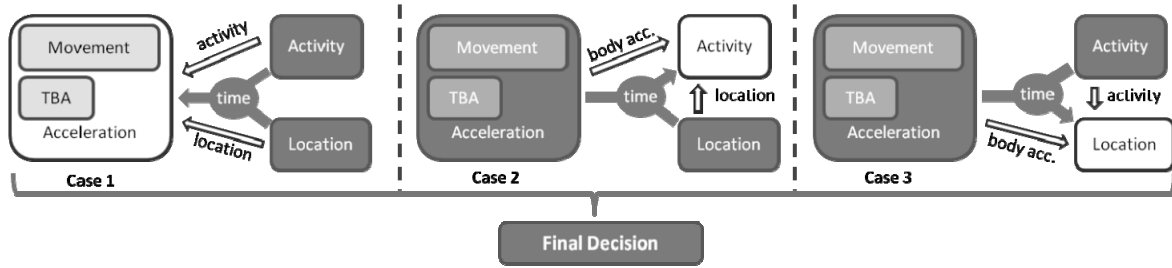


Fig. 5. Context-based reasoning schema.

Table 1

Context variables and their possible values. *A* – activity, *L* – location, *B_M* – body movement

	<i>A</i>	<i>L</i>	<i>B_M</i>
standing	sitting unusual	bed	yes
sitting	on all fours	floor	no
lying	going down	chair	
	standing up		

The reasoning about the user’s situation was performed using expert rules. An example rule that describes a situation when the person is lying (“Ly”) on the floor (“F”) without movement in the time interval $[t1, t2]$ is shown in (3):

$$(A[t1,t2] = \text{“Ly”}) \parallel (B_M[t1,t2] = \text{“no”} \wedge L[t1,t2] = \text{“F”}) \quad (3)$$

The rule consists of two parts: the base (left-side) and the contextual (right-side) part, separated by \parallel . The base part represents the trigger which signals when the contextual part should be checked. For the example rule given in (3), first the lying activity is recognized in the interval $[t1, t2]$. This triggers the analysis of the context variables: body movement (*B_M*) and location (*L*). We selected the time interval *t* for the reasoning about the user’s situation to be 10 seconds. This way, the interval is long enough for a reliable recognition, but still negligible compared to the time needed for help to arrive.

During the reasoning time interval, a context variable may contain different values, e.g., 5 lying, 3 standing up and 2 sitting activities. In order to represent the whole interval with one context value, we empirically selected 80% as the minimum percentage of same values that a context variable should have (e.g., the activity should be lying 8 seconds out of 10 to satisfy a rule that requires lying). Otherwise, if the context values are represented with smaller percentage than 80%, none of the context values is chosen for the particular interval.

Given the values of the context variables shown in Table 1, there are 12 possible cases that can be used

as a trigger to a rule. Each of these cases can be evaluated in the context of a one or both other context components, resulting in 249 possible rules. However, because in this study we only focused on fall situations, it turned out that only a few context-based rules cover all the falls analyzed. Once we established the general reasoning scheme, adding more rules for different situations is a relatively easy task.

8.1. Inertial + location sensors reasoning

The reasoning rules, when the two types of sensors were combined, primarily relied on the recognized activity; the additional context information consisted of the location and the body movement. A fall situation was defined by each of the following rules:

$$(A[t1, t2] = \text{“Ly”}) \parallel (B_M[t1, t2] = \text{“no”} \wedge L[t1, t2] = \text{“F”}) \quad (4)$$

$$(A[t1, t2] = \text{“SU”}) \parallel (B_M[t1, t2] = \text{“no”} \wedge L[t1, t2] = \text{“F”}) \quad (5)$$

$$(A[t1, t2] = \text{“OA4”}) \parallel (B_M[t1, t2] = \text{“no”} \wedge L[t1, t2] = \text{“F”}) \quad (6)$$

“Ly” stands for the activity *lying*, “SU” for *sitting unusual*, “OA4” for *on all fours*, and “F” stands for the *floor* location.

We used assumptions that the elderly do not usually lie or sit on the ground and are not on all fours for more than *t* seconds while not moving.

8.2. Inertial sensors reasoning

In this subsection we present the context-based reasoning when a conclusion is inferred based on inertial sensors alone and therefore a fall situation is defined using the activity and the body accelerations context variables.

Our previous experiments and also some related work [19,33], showed that it was possible to detect a straightforward (fast) fall by using only TBA; however, lots of false positives appeared in other fall-like events: quickly lying down on a bed, quickly sitting

Table 2
The events in the scenario, the appropriate activities and event descriptions

	#	Event	Activities	Description
Fall Events	(1)	Fast fall (tripping)	Standing/walking, going down, lying, standing up	Falling performed in different ways: forwards, backwards or to the sides
	(2)	Slow fall (fainting)	Standing/walking, going down, lying, standing up	Losing consciousness and slowly falling to the ground (trying to hold onto furniture)
	(3)	Falling when trying to stand up	Sitting, standing up, going down, sitting on the ground, standing up	Trying to stand up from a chair, but having difficulties and slowly falling to the ground, ending up in a sitting posture on the ground
	(4)	Sliding from a chair	Sitting, standing up, going down, sitting on the ground	Sliding from a chair and ending up in sitting unusual on the ground
Fall-like Events	(5)	Quickly lying down on a bed	Standing/walking, going down, lying, standing up	Quickly lying down on a bed
	(6)	Quickly sitting down on a chair	Standing/walking, going down, sitting, standing up	Quickly sitting down on a chair
	(7)	Searching for something on the ground	Standing/walking, going down, on all fours, lying	Going on all fours and afterwards going to lying posture in order to take an object from the ground
Normal Events	(8)	Sitting down	Standing/walking, going down, sitting, standing up	Sitting down on a chair normally
	(9)	Lying down	Standing/walking, going down, lying, standing up	Lying down on a bed normally
	(10)	Walking	Standing/walking	Walking sequences between events

on a chair, etc. Therefore, a potential fall detected by TBA was confirmed by the body movement and additional context information, i.e., the user's activity. As an example, a fall situation is defined by each of the following rules:

$$(TBA[t1] = \text{"yes"}) \parallel (A[t1, t2] = \text{"Ly"} \wedge B_M[t1, t2] = \text{"no"}) \quad (7)$$

$$(A[t1, t2] = \text{"SU"}) \parallel (B_M[t1, t2] = \text{"no"}) \quad (8)$$

$$(A[t1, t2] = \text{"OA4"}) \parallel (B_M[t1, t2] = \text{"no"}) \quad (9)$$

$TBA[t1] = \text{"yes"}$, represents the time when a high-acceleration fall pattern is detected.

8.3. Location sensors reasoning

Since the location sensors are better at AR than detecting fall accelerations, FD is based on the activity that may result from a fall, and uses location as the context. The first advantage compared to the stand-alone inertial FD was the location information: the system was aware of some predefined "safe" locations, such as the bed. The second advantage was the z coordinate of the sensor location, which provides the height of the sensor and therefore distinguishes different activities, for example, sitting on the floor from sitting on a chair. An example of a rule structure is presented here:

$$(A[t1, t2] = \{\text{"Ly"} \vee \text{"SU"} \vee \text{"OA4"}\}) \parallel L[t1, t2] = \text{"F"} \quad (10)$$

9. Experimental setup

9.1. Experimental scenario

A complex, 15-minute test scenario was specifically designed to investigate events that might be difficult to recognize as falls or non-falls. This scenario, shown in Table 2, was created in consultation with a medical expert. In Table 2 the numbers in parentheses represent the event numbers for easier referencing throughout the text. The events were recorded in a single recording including all the events.

Because typical fast falls are easy to detect due to high acceleration, only one such fall (1) was included. Three atypical falls not involving high acceleration, i.e., (2), (3) and (4), were included to test the use of the contextual activity information, i.e., that a person is not expected to sit/lay on the ground (as opposed to the chair/bed). Furthermore, the two events (5) and (6) involve high acceleration and could thus be misclassified as falls by acceleration-based methods, such as TBA. However, the methods that use the activity and location as contextual information should be able to detect that these are non-fall events. An event (7) was included that involves voluntarily lying on the ground, which could mislead the methods that use information other than acceleration. The events (8), (9) and (10) are normal and were included to

verify that all the methods work correctly during normal events.

Additionally all the target activities mentioned in Section 7.2 are contained in the scenario (see Table 2): standing, sitting, lying, sitting unusual, on all fours, going down and standing up.

The experimental scenario was recorded with all 6 inertial and 4 location sensors. Afterwards, the CoFDILS was tested with all 1023 combinations of sensors (single type, as well as both types).

The scenario was recorded by 11 young healthy volunteers (24–33 years, 7 males and 4 females). It was repeated 5 times by each person, resulting in 55 recordings and a total of 550 events for the FD and total number of 105 438 segmented samples for the AR. Testing elderly people was not feasible because the scenario was too strenuous and risky for them, but the volunteers were advised how to act by the medical expert in order to mimic elderly. Additionally, the data for 3 more people was recorded for tuning the basic parameters, e.g., thresholds, preliminary tests and choosing the best algorithms.

9.2. Evaluation metrics

For the evaluation of the AR, the leave-one-person-out cross-validation technique was used. This means the model was trained on the data recorded for ten people and tested on the remaining person. This procedure was repeated for each person (11 times). This evaluation approach is more reliable than using the same persons' data for training and testing. Using the same person's data would give overly optimistic results if the intended use of the model is to classify the activities of previously unseen people. Several evaluation metrics were analyzed, such as the recall, precision, accuracy and F-measure. For the final conclusions, the F-measure was chosen, since it is a harmonic mean of both precision and recall [31] and thus treats false negatives and false positives in a balanced way. First, for each fold the weighted averaged F-measure (over all activities) was calculated and finally the average F-measure over all the folds was calculated.

To evaluate the FD, one must decide how to weigh the undetected falls and the false alarms. Both are important: not detecting a fall may endanger a person's health, while false alarms make the system unlikely to be used in real life. Therefore, we used F-measure (F), which weights undetected falls and false alarms equally. It is defined as a harmonic mean of recall (the percentage of the events recognized as

falls/non-falls from all the fall/non-fall events) and precision (the percentage of the events truly being falls/non-falls of all the events recognized as such) [31]. In the detailed results in Table 8, the true positive and true negative rates are also presented.

10. Experimental results and discussion

The experimental results and discussions are provided in the following two subsections. Tests to confirm the statistical significance of the results were also performed. The best sensor combinations for each number of sensors and each sensor type were tested separately. Because of the small number of folds (11) and because the individual samples are paired (the same person's data for each combination), we used paired Student's T-test with a significance level of 5%.

10.1. Activity recognition

Figure 6 presents a detailed matrix (5×7) of the best sensor combinations for the AR. The inertial sensors are shown on the horizontal axis and the location on the vertical axis. Each rectangle in the matrix contains the sensors placements and the achieved F-measure marked with F in percent. For example, the (2, 3) rectangle represents the best combination of 2 location and 3 inertial sensors. It is also the best of all sensor combinations according to the F-measure = 98.6%. The dotted lines (diagonal) connect the rectangles that have the same number of sensors. Along each dotted line the best performing sensor combination is marked with a black circle.

Another representation of the same results is shown in Fig. 7. This is a 3D representation, where the third axis is the performance – F-measure. The points of the graph that are closer to 0 in the xy plane (smaller number of sensors), are of the greatest interest for practical usage. These points are marked with circles (one-sensor type) and squares (combined-sensor type).

Analyzing the results achieved by the inertial sensors only (horizontal axis in Fig. 6), one can see that the performance improves by adding inertial sensors up to three (see Fig. 7). This was also proved by the significance test. Adding more sensors does not significantly increase the performance; it only makes the system more intrusive and complex.

The analysis of the AR results for the location sensors (vertical axis in Fig. 6) is different. The increase

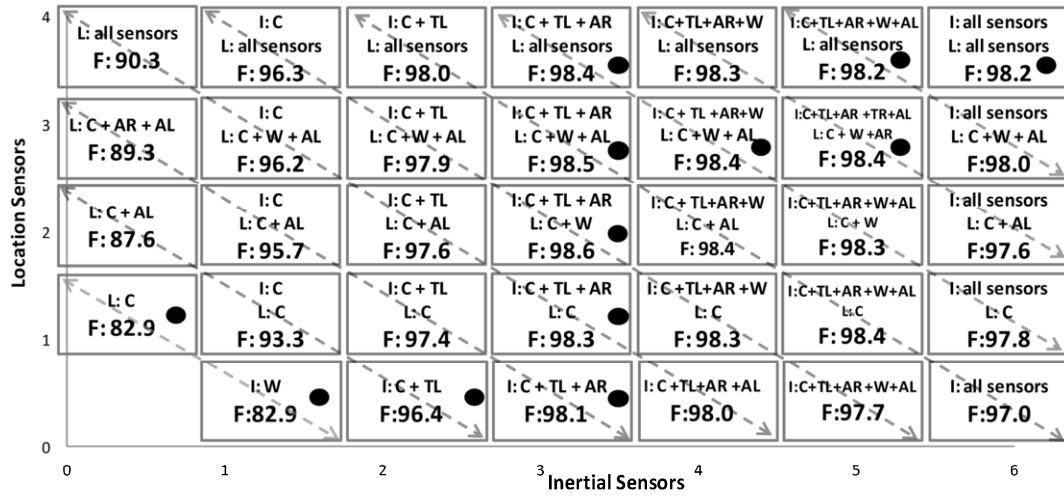


Fig. 6. Activity recognition detailed results. Matrix representation of the best sensor combinations using the Inertial (I) and Location (L) sensors. F – overall F-measure, C – Chest, W – Waist, AR – Ankle Right, AL – Ankle Left, TR – Thigh Right, TL – Thigh Left.

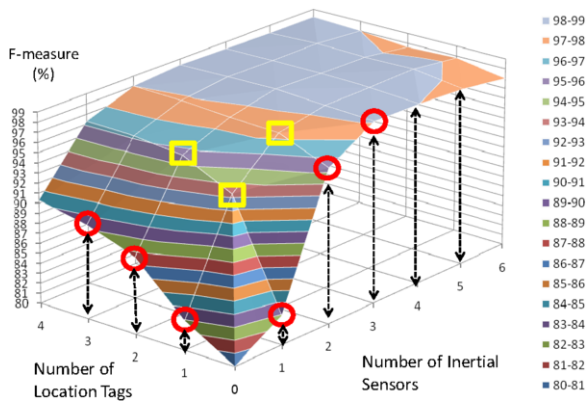


Fig. 7. Best sensor combinations for activity recognition.

in the number of sensors has a direct influence on the increase in the system performance. There is no convergence or decrease in performance as sensors are added, as is the case with the inertial sensors only. The statistical tests additionally confirmed this.

The results for the combined sensors showed that the difference in performance is statistically significant only when the system is using one, two and three sensors. Four sensors or more do not significantly increase the performance of the system. In general, the results showed that for achieving the highest performance only three sensors are needed.

The improvements of the combined sensors compared to the location sensors alone are evident (Fig. 7 squares vs. circles). However, when the combined sensors are compared to the inertial sensors alone, the

Table 3

Activity recognition analysis using only 1 sensor enclosure equipped with inertial and location sensor

1-sensor enclosure (Inertial + Location)			
C	W	AL	AR
93.3%	90.9%	85.0%	84.4%

Table 4

Activity recognition analysis using only 2 sensor enclosures equipped with inertial and location sensor

2-sensor enclosures (Inertial + Location)					
C+AL	C+AR	C+W	W+AR	W+AL	AR+AL
94.5%	93.6%	92.2%	92.1%	91.5%	84.7%

situation is different. With two and three sensors, the inertial sensors alone outperformed the combination of both sensors. This means that for the AR, it is better to add an inertial than a location sensor. This may be a bit surprising, but the probable explanation is that the inertial sensors are less noisy and provide sensor inclination angles, which are well-suited to activity recognition, especially for the static activities.

Since both sensor types can be put in the same enclosure, one can also examine the number of enclosures. Tables 3 and 4 show the results for the AR when 1 and 2 sensor enclosures (equipped with both sensor types) are analyzed, respectively.

The results in Fig. 6 show that the combination of 1 inertial and 1 location sensor on the chest clearly outperforms each of the sensors used separately. The achieved F-measures are 83% for each of the systems used separately. Their combination improves these

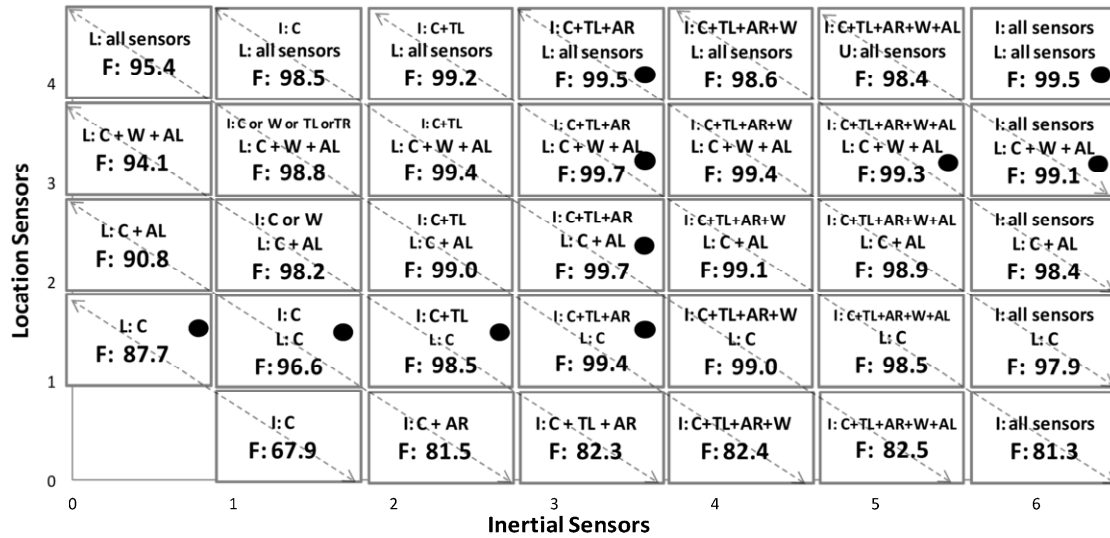


Fig. 8. Matrix representation of the best sensor combinations using the Inertial (I) and Location (L) sensors. F – overall F-measure, C – Chest, W – Waist, AR – Ankle Right, AL – Ankle Left, TR – Thigh Right, TL – Thigh Left.

results significantly (i.e. by 10 pp.). Furthermore, the results in Table 3 show that the combination of 1 inertial and 1 location sensor placed on the chest outperforms each of the other sensor placement combinations: waist, left and right ankle.

The results in Table 4 show that the best performing 2-enclosures-placement (when both types of sensors are included in each enclosure) is the chest and left ankle achieving 94.5% performance. However, this is not the best performing 2-enclosures-placement, because the combination of 1 inertial and 1 location sensor on the chest and 1 inertial sensor in the left thigh (shown in Fig. 6) achieves 97.4% performance. Please note that thigh placements were not tested for the location sensors. This shows that it is better to add inertial sensor alone on the thigh (97.4%) instead of adding the both (inertial and location) to the left ankle (96.5%). The reason for this is that the inertial sensors are better for AR and that the thigh placement is better compared to the ankle.

The best performing combination of three sensor enclosures is chest (inertial and location), right ankle (inertial and location) and left thigh (inertial only).

10.2. Fall detection

Similar to the AR, Fig. 8 presents a matrix (5 × 7) of the best sensor combinations for the FD. The inertial sensors are shown on the horizontal axis and the location on the vertical axis. Each rectangle in the matrix contains the sensor placements and the achieved F-measure marked with F as a percentage.

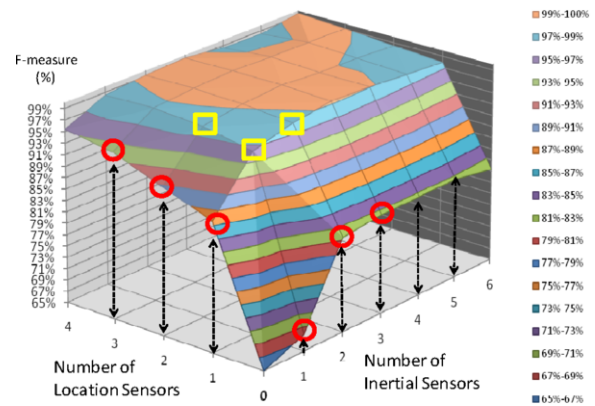


Fig. 9. Best sensor combinations for fall detection.

For example, the (2, 3) rectangle represents the combination of 2 location and 3 inertial sensors. It is the best of all combinations according to the F-measure = 99.7%. The dotted lines (diagonal) connect the rectangles that have the same number of sensors. Along each dotted line the best (according to the F-measure) rectangle is marked with a black circle. These rectangles represent the best combination given the number of sensors.

Another representation of the same results is shown in Fig. 9. This is a 3D representation, where the third axis is the achieved F-measure.

Analyzing the results achieved with the inertial sensors alone (Fig. 8 horizontal axis rectangles), one can see that the only important improvement is achieved when using two sensors instead of one. Af-

Table 5

Fall detection analysis using only 1 sensor enclosure equipped with inertial and location sensor

1-sensor enclosure (Inertial + Location)			
C	W	AL	AR
96.1%	95.6%	75.0%	72.1%

Table 6

Fall detection analysis using only 2 sensor enclosures equipped with inertial and location sensor

2-sensor enclosures (Inertial + Location)					
C+AL	C+AR	C+W	W+AR	W+AL	AR+AL
98.3%	98.0%	97.2%	97.6%	97.9%	83.7%

ter this, adding up to five sensors did not significantly improve the F-measure; including a sixth sensor even decreased the performance.

For the location sensors, an increase in the number of sensors increases the performance all the way. The statistical tests proved that there is a significant difference in the performance of the system using one, two, three and four location sensors. Like with the inertial FD, the chest is the best-performing placement.

The statistical tests for the combined FD showed that the difference in performance is statistically significant only when the system is using two and three sensors. Four sensors or more do not significantly increase the performance of the system.

The parts of the graph with a smaller number of sensors are of the greatest interest for practical usage (squares and circles in Fig. 9). The combination of sensors clearly outperforms the individual sensor types. For example, the performance values of the system using two sensors are 81.5% and 90.8%, for the inertial and location sensors, respectively. Their combination achieves 96.6%, an improvement of 15 pp. and 6 pp., respectively. This is the case for each number of sensors (dotted lines): the combination of two sensor types is better than each of the types used separately.

Since both sensor types can be put in the same enclosure, one can also examine the number of enclosures. Tables 5 and 6 show the results for the FD when 1 and 2 sensor enclosures (equipped with the both sensor types) are analyzed, respectively.

The performance of the system using only one sensor of one type is 68% and 88% for the inertial and location sensor, respectively. The results in Table 5 show that by combining them into one enclosure on the chest, the achieved F-measure is 96.6%, an im-

Table 7

The simplest and the best (statistically significant) combinations of the inertial-only, location-only and both types of sensors

	The simplest combination	The best combination
Inertial sensors	Chest	Chest + Right ankle
Location sensors	Chest	All four sensors
Combined sensors	Inertial: Chest Location: Chest	Inertial: Chest + Right ankle Location: Chest

provement of 29 pp. and 9 pp., respectively. Furthermore, the combination of 1 inertial and 1 location sensor placed on the chest outperforms each of the other sensor placement combinations: waist, left and right ankle.

The results in Table 6 show that the best performing 2-enclosures-placement (when both types of sensors are included in each enclosure) is the chest and left ankle achieving 98.3% performance. However, this is not the best performing 2-enclosures-placement, because the combination of 1 inertial and 1 location sensor on the chest and 1 inertial sensor in the left thigh (shown in Fig. 8) achieves 98.5% performance. This shows that it is better to add inertial sensor alone on the thigh (98.5%) instead of adding the both (inertial and location) to the left ankle (98.3%). The reason for this is the improvement in the AR module, which is greater when the thigh inertial sensor is introduced compared to the ankle inertial and location. The best performing combination of three sensor enclosures is chest (inertial and location), right ankle (inertial and location) and left thigh (inertial only).

The rest of the discussion is a detailed analysis of the results achieved by the statistically significant simplest and the best combinations of the inertial-only, location-only and both types of sensors. The sensor types and placements are shown in Table 7 and the results are presented in Table 8. The events in Table 8 are divided into three groups: fall, non-fall (fall-like), and normal events. The numbers are the percentage of all fall/non-fall events being correctly recognized as fall/non-fall (true positive and negative rate). The last row represents the overall F-measure.

The first two columns show the results achieved for the FD with inertial sensors. The first event in Table 8, tripping, is a typical fall that was recognized accurately because of the TBA rule. The second event, which is falling slowly, was difficult to recognize because of the low acceleration during this event.

Table 8
Detailed FD results for each event and each context-based FD method

		Context-based Reasoning					
		Inertial (Activity + TBA + Movement)		Location (Activity + Location)		Combination (Activity + TBA + Movement + Location)	
		Simplest	Best	Simplest	Best	Simplest	Best
Fall Events	(1) Tripping – Quick falling	100	100	96	100	100	100
	(2) Fainting – Falling slowly	11	11	100	100	100	100
	(3) Falling from a chair slowly	68	98	95	95	99	99
	(4) Sliding from a chair	72	99	97	97	98	99
Non-Fall Fall-like Events	(5) Sit down quickly on a chair	55	97	75	89	91	98
	(6) Searching on the ground	85	88	25	78	80	89
	(7) Quickly lying down on a bed	34	34	100	100	100	100
Non-Fall Normal Events	(8) Sitting normally	68	98	80	93	93	98
	(9) Lying normally	100	100	100	100	100	100
	(10) Walking	97	100	92	97	100	100
Overall F-measure in %		67.9	81.5	87.7	95.4	96.6	98.5

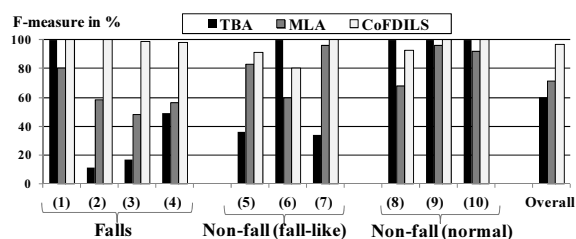


Fig. 10. Comparison of the FD results achieved by our Context-based approach, the Machine-learning approach (MLA), and Threshold-based approach (TBA). The event numbers correspond to the events given in Table 2.

For this event, additional contextual information was necessary (e.g., the location of the user). The effect of the activity information of the user can be seen in the fall events that end with sitting unusual on the ground (events 3 and 4). In these cases the AR model correctly recognized sitting unusual on the ground. On the other hand, this has a negative impact on the performance when the sitting event is analyzed (events 5 and 8). In this case, the AR model was not accurate enough and recognized sitting unusual on the ground, resulting in a false positive. This issue was solved by including more sensors, which improved the AR method (e.g., the column Inertial-best).

The location sensors based FD was using the activity and the location information. Because of the location, it recognized all falls with high accuracy

(events 1 to 4). However, some problems remained among the non-fall events, because of the relatively low accuracy of the AR model. Namely, sitting (events 5 and 8) and searching on the ground (event 6) were misclassified as sitting unusual on the ground or lying (on the ground), causing the system to detect a fall during the non-fall events. Improvements in the performance can be seen when the number of sensors is increased (the column Location-best), due to the improvements in the AR method.

The last two columns show the results achieved with the combination of both types of sensors and the full context as presented in Section 7.1. The improvements are clear in all of the events. The overall performance when two sensors (one inertial and one location) were used was 96.6%. Some problems only appeared among the non-fall events that ended with sitting (5 and 8) and the searching on the ground event (6). The reason lies in the AR method, which misrecognized the appropriate activities (sitting and on all fours). These problems were solved by including one more inertial sensor, which significantly improves the AR model and consequently the FD (the last column in Table 8).

Finally, two commonly used methods in the literature, the threshold-based approach (TBA) and the machine-learning approach (MLA), were tested for comparison. The results are shown in Fig. 10, by presenting the true-positive or true-negative rate for each fall or non-fall event, respectively. The TBA is

described in Section 7.1.1. More details about the MLA can be found in our previous work (Luštrek et al. [20]). The basic principle of MLA is that a machine-learning model is trained to detect a fall event. In our case, features extracted from the chest-inertial and chest-location sensor data were used. Therefore, the contextual location information was implicitly (through features) introduced in the MLA.

The overall results showed that our method, in which the context is explicitly encoded with rules, outperformed the other two methods, which use: implicit context information (MLA) or only accelerations (TBA). The TBA outperformed our CoFDILS only in two events (6, 8); however, this was due to the one-sided performance of the TBA (detects only high accelerations) at the expense of the overall performance.

11. Conclusion

We presented a novel approach for fall detection called CoFDILS, which combines inertial and location sensors using a general context-based schema. The method exploits three context components to detect a fall situation: the activity of the user, the body accelerations and the location. Each component is formally presented with a context variable that contains the value of the component at each point in time. The reasoning over the values of the context variables is performed by expert rules. There are 249 possible combinations of the values of the context variables (rules); however, in our tests it turned out that only a few of the context-based relations capture most of the context information in the FD domain. Once we established the general reasoning scheme, adding more rules for different situations is a relatively trivial task. Also the addition of new sensors, such as sound and vibration, is relatively easy from the reasoning point of view; just adding a context variable and including it in the rules. Currently, the context-based reasoning rules were designed manually. The automation of learning the best context relations is considered for future work.

We tested the performance with all possible combinations of the six inertial and four location sensors to find the best sensor placements, using the context-based reasoning schema. The evaluation was performed on a complex test scenario, which included real-life, realistic events that are difficult to recognize as falls or non-falls. The results showed that by combining the two types of sensors it is possible to detect complex fall situations by using the activity and the

context information from both types of sensors. It is essential that both sensor types are employed, since they provide complementary information about the user's situation. Finally, the best practical solution proved to be the chest placement with a single sensor enclosure including one inertial and one location sensor achieving 96.6% for the fall detection employed in CoFDILS and 93.3% for the activity recognition task only.

For the future work, we plan to focus more on the practical implementation of the system. First, we intend to make the system less intrusive and more user-friendly. This can be achieved by using completely wireless inertial sensors. Another improvement in this direction can be achieved by introducing specially designed clothes, which will include pockets for the sensors. Additionally, the interaction between the user and the system can be improved. This can be achieved by including the user's smartphone, tablet or PC as a medium for showing system's notifications (fall detected, system malfunction, etc.).

Acknowledgments

This work was partly supported by the Slovenian Research Agency under the Research Programme P2-0209 Artificial Intelligence and Intelligent Systems and partly from the European Community's Framework Programme FP7/2007-2013 under grant agreement No. 214986. The authors would like to thank Boštjan Kaluža for the helpful suggestions and discussions.

References

- [1] J.A. Botía, A. Villa, J.T. Palma, D. Pérez and E. Iborra, Detecting domestic problems of elderly people: simple and unobtrusive sensors to generate the context of the attended, in: *Proc. of the 1st International Workshop on Ambient Assisted Living*, Vol. 5518, 2009, pp. 819–826.
- [2] L. Breiman, Random forests, in: *Machine Learning*, Vol. 45, 2001, pp. 5–32.
- [3] A. Dey, D. Salber, G. Abowd and M. Futakawa, The conference assistant: combining contextcon-text-awareness with wearable computing, in: *Proc. in International Symposium on Wearable Computers*, 1999.
- [4] C. Doukas and I. Maglogiannis, Advanced patient or elder fall detection based on movement and sound data, in: *Proc. Pervasive Computing Technologies for Healthcare*, 2008, pp. 103–107.
- [5] C.N. Doukas and I. Maglogiannis, Emergency fall incidents detection in assisted living environments utilizing motion, sound, and visual perceptual components, *IEEE Transactions on Information Technology in Biomedicine* **15**(2) (2011), 277–289.

- [6] S.M. Friedman, B. Munoz, S.K. West, G.S. Rubin and L.P. Fried, Falls and fear of falling: which comes first? A longitudinal prediction model suggests strategies for primary and secondary prevention, *Journal of the American Geriatrics Society* (2002), 1329–1335.
- [7] H. Gjoreski, M. Gams and I. Chorbev, 3-axial accelerometers activity recognition, in: *Proc. ICT Innovations*, 2010, pp. 51–58.
- [8] H. Gjoreski, M. Luštrek and M. Gams, Accelerometer placement for posture recognition and fall detection, in: *Proc. the 7th International Conference on Intelligent Environments*, 2011, pp. 47–54.
- [9] H. Gjoreski, M. Luštrek and M. Gams, Context-based fall detection using inertial and location sensors, in: *Proc. International Joint Conference on Ambient Intelligence*, Lecture Notes in Computer Science, Vol. 7683, 2012, pp. 1–16.
- [10] L.D. Gillespie, M.C. Robertson, W.J. Gillespie, S.E. Lamb, S. Gates, R.G. Cumming and B.H. Rowe, Interventions for preventing falls in older people living in the community (Review), *The Cochrane Library*, 4, 2009.
- [11] D. Gimón, H. Gjoreski, B. Kaluža and M. Gams, Using accelerometers to improve position-based activity recognition, in: *Proc. 13th International Multi-Conference Information Society*, 2010, pp. 15–18.
- [12] P. Jantaraprim, P. Phukpattaranont, C. Limsakul and B. Wongkittisuksa, Evaluation of fall detection for the elderly on a variety of subject groups, in: *Proc. The 3rd International Convention on Rehabilitation Engineering & Assistive Technology*, 2009.
- [13] B. Kaluža and E. Dovgan, Glajenje trajektorij gibanja cloveskega telesa zajetih z radijsko tehnologijo, in: *Proc. International Multi-Conference Information Society*, 2009.
- [14] S. Khawandi, B. Daya and P. Chauvet, Automated monitoring system for fall detection in the elderly, *Proc. International Journal of Image Processing (IJIP)* 4(5) (2010).
- [15] S. Kozina, H. Gjoreski, M. Gams and M. Luštrek, Three-layer activity recognition combining domain knowledge and meta-classification, *Journal of Medical and Biological Engineering* 33(4) (2013), 406–414.
- [16] S. Kozina, H. Gjoreski, M. Gams and M. Luštrek, Efficient activity recognition and fall detection using accelerometers, in: *Evaluating AAL Systems Through Competitive Benchmarking*, Communications in Computer and Information Science, Vol. 386, 2013, pp. 13–23.
- [17] Q. Li, J. Stankovic, M. Hanson, A.T. Barth, J. Lach and G. Zhou, Accurate, fast fall detection using gyroscopes and accelerometer-derived posture information, in: *Proc. Sixth International Workshop on Wearable and Implantable Body Sensor Networks*, 2009, pp. 138–143.
- [18] Q. Li and J. Stankovic, Grammar-based, posture and context-cognitive detection for falls with different activity levels, in: *Proc. WH'11: 2nd Conference on Wireless Health*, 2011.
- [19] H.T. Lin, Y.J. Hsieh, M.C. Chen and W. Chang, ActionView: a movement-analysis ambulatory monitor in elderly homecare systems, in: *IEEE International Symposium on Circuits and Systems*, 2009, pp. 3098–3101.
- [20] M. Luštrek, H. Gjoreski, S. Kozina, B. Cvetkovic, V. Mirchevska and M. Gams, Detecting falls with location sensors and accelerometers, in: *Proc. Twenty Third IAAI Conference*, 2011.
- [21] J. Machajdik, S. Zambanini and M. Kampel, Fusion of data from multiple cameras for fall detection, in: *Proc. Workshop on Behaviour Monitoring and Interpretation, BMI'10*, 2010, pp. 1–7.
- [22] A. Muñoz, E. Serrano, A. Villa, M. Valdés and J.A. Botía, An approach for representing sensor data to validate alerts in ambient assisted living, *Sensors* 12 (2012), 6282–6306.
- [23] H.D. Naranjo, L.M. Roa, T.J. Reina and V.M.A. Estudillo, Personalization and adaptation to the medium and context in a fall detection system, *IEEE Trans. Inf. Technol. Biomed.* 16(2) (2012), 264–271.
- [24] T.T. Nguyen, M.C. Cho and T.S. Lee, Automatic fall detection using wearable biomedical signal measurement terminal, in: *Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 2009, pp. 5203–5206.
- [25] L.Z. Rubenstein and K.R. Josephson, The epidemiology of falls and syncope, *Clinics in Geriatric Medicine* 18(2) (2002), 141–158.
- [26] A. Sixsmith and N.A. Johnson, Smart sensor to detect the falls of the elderly, *Pervasive Computing, IEEE* 3(2) (2004), 42–47.
- [27] E.M. Tapia, Using machine learning for real-time activity recognition and estimation of energy expenditure, PhD thesis, Massachusetts Institute of Technology, 2008.
- [28] M.E. Tinetti and C.S. Williams, Falls, injuries due to falls, and the risk of admission to a nursing home, *The New England Journal of Medicine* 337 (1997), 1279–1284.
- [29] Ubisense RTLS website, <http://www.ubisense.net> (accessed February 2013).
- [30] World Health Organization (WHO), Global brief for World Health Day 2012, Good health adds life to years, http://whqlibdoc.who.int/hq/2012/WHO_DCO_WHD_2012_2_eng.pdf (accessed February 2013).
- [31] I. Witten and E. Frank, *Data Mining: Practical Machine Learning Tools and Techniques*, 2nd edn, Morgan Kaufmann, 2005.
- [32] XSens sensors website, <http://www.xsens.com> (accessed February 2013).
- [33] T. Zhang, J. Wang, L. Xu and P. Liu, Fall detection by wearable sensor and one-class SVM, *Intelligent Computing in Signal Processing and Pattern Recognition* 345 (2006), 858–863.
- [34] A. Zinnen, C. Wojek and B. Schiele, Multi activity recognition based on bodymodel-derived primitives, in: *Proc. 4th Int. Symposium on Location and Context Awareness*, 2009.