Fall Detection Using Location Sensors and Accelerometers

Version 3.1, submitted 2014-11-28

Mitja Luštrek[1,2], Hristijan Gjoreski[1,2], Narciso González-Vega[3], Simon Kozina[1], Božidara Cvetković[1,2], Violeta Mirchevska[1,4], Matjaž Gams[1,2]

[1] Jožef Stefan Institute, Department of Intelligent Systems, Jamova cesta 39, 1000 Ljubljana, Slovenia
[2] Jožef Stefan International Postgraduate School, Jamova cesta 39, 1000 Ljubljana, Slovenia
[3] University of Jyväskylä, Agora Human Technology Center, Mattilanniemi 2, 40014 Jyväskylä, Finland
[4] Result d. o. o., Bravničarjeva 11, 1000 Ljubljana, Slovenia

{mitja.lustrek, hristijan.gjoreski, simon.kozina, boza.cvetkovic, violeta.mircevska, matjaz.gams}@ijs.si
narciso.gonzalez-vega@jyu.fi

Abstract:

The rapid aging of the population drives the development of pervasive solutions for the care of the elderly, which often involve fall detection with accelerometers. These solutions are very accurate in laboratory conditions, but can fail in some real-life situations. To overcome this, we present the Confidence system, which detects falls mainly with location sensors. A user wears one to four tags whose locations are detected with sensors. This allows recognizing the user's activity, including falling and lying afterwards, and the context in terms of the location in the apartment. A scenario consisting of events difficult to recognize as falls or non-falls was used to compare the Confidence system with accelerometer-based fall-detection methods, some of them augmented with the context from a location sensor. The accuracy of the methods that utilized the context was around 30 percentage points higher compared to the methods without the context. The Confidence system was also successfully validated in a real-life setting with elderly users.

Keywords:

I.2 Artificial Intelligence, J.3.b Health applications, J.9.e Wearable computers and body area networks

1 Introduction

The world population is aging rapidly, threatening to overwhelm the society's capacity for taking care of its elderly members. In the developed world, there were four persons aged 15–64 for every person aged

65+ in 2012, but the United Nations project that the ratio will decline to two-to-one by 2050. This drives the development of technical solutions to help the elderly live longer independently.

There are at least two reasons why fall detection is a very active research topic. First, falls are a serious problem: approximately half of the hospitalizations of the elderly are caused by falls, the fear of falling is an important cause for nursing home admission, and "the long lie" (not being able to get up and call for help) is a good predictor of death within six months[1]. And second, falls can be detected fairly effectively with the currently available pervasive technology. This is most often done with accelerometers, which can measure the high acceleration upon the impact with the ground. Accelerometers are accurate, lightweight and inexpensive; however, their limitations are that some safe activities result in high acceleration, and more importantly, not all falls result in high acceleration.

We developed an alternative approach to fall detection in the Confidence project (http://www.confidence-eu.org). A Confidence user wears one to four tags on the body, whose locations are detected with sensors. From the tag locations the user's activity and location in the apartment are inferred. This can be augmented with acceleration data from an accelerometer, resulting in rich information enabling reliable detection even of atypical falls. The main motivation for our approach is that the same activity should be treated as a fall or not depending on the context – e.g., lying in bed is practically always a safe activity, while lying on the ground is very often related to a fall, especially in the elderly. Other contextual information can be taken into account, such as whether the user is moving. A video that shows an overview of our approach and motivation can be seen here: http://www.youtube.com/watch?v=1euk8G3pEoM.

2 Related Work on Fall Detection

Most research on fall detection uses accelerometers and sometimes gyroscopes. Falls are detected either by applying thresholds to accelerations, velocities and angles[2], or by using machine learning to recognize the signal pattern corresponding to a fall[3,4]. In addition to recognizing the fall itself, a person's posture after a potential fall is often detected, since most falls result in the person lying on the ground. The accuracies reported by the authors of fall-detection methods are very high, sometimes 100%, making the impression that fall detection is largely a solved problem. However, these approaches were mostly tested on typical fast falls simulated by young volunteers and common daily activities. A test on real-life falls by the elderly showed markedly worse accuracy[2].

Of particular interest to us is the work by Li et al.[5], where fall detection was tested on simulated but more difficult cases. Their average accuracy was 90.1% because of the poor performance on two atypical fall types and lying down quickly. The authors later addressed this problem by utilizing the context, which consisted of furniture usage measured with pressure sensors and additional information on the user's movement[6]. This resulted in the accuracy of 97.8% on similar test cases as in the earlier work.

We also addressed the problem of atypical falls and daily activities difficult to distinguish from falls by using context. However, in our case it consisted of the user's location measured with location sensors[7,8]. In this article we expand upon our earlier work by including additional fall-detection methods and presenting results of experiments on elderly volunteers.

Fall detection using location sensors is quite rare, probably because common location-sensing technologies, such as wireless LAN or Bluetooth, offer accuracy on the order of 1 m or worse[9], making them unsuitable for detecting falls on their own. Besides our work, Bowen et al.[10] detected falls by simply determining whether an ultra-wideband radio tag attached to the wrist is located on the ground or not, which achieved the accuracy of 80.5–89.0% on clear-cut cases.

3 Fall Detection Methods

We detected falls with the Confidence system, which mainly uses location sensors but can augment them with an accelerometer, and with accelerometer-based methods, which can be augmented with location information.

3.1 The Confidence System

Figure 1 shows the fall-detection pipeline in the Confidence system. The data from location sensors is first preprocessed and then used for activity recognition. Afterwards, the recognized activities are combined with the context in terms of the user's location to detect falls, which may trigger an alarm. Optionally, an accelerometer may be used to correct the activity recognition and augment the context used in the fall detection (dotted in the figure). The Confidence system also adapts to each user (dashed/dotted in the figure).

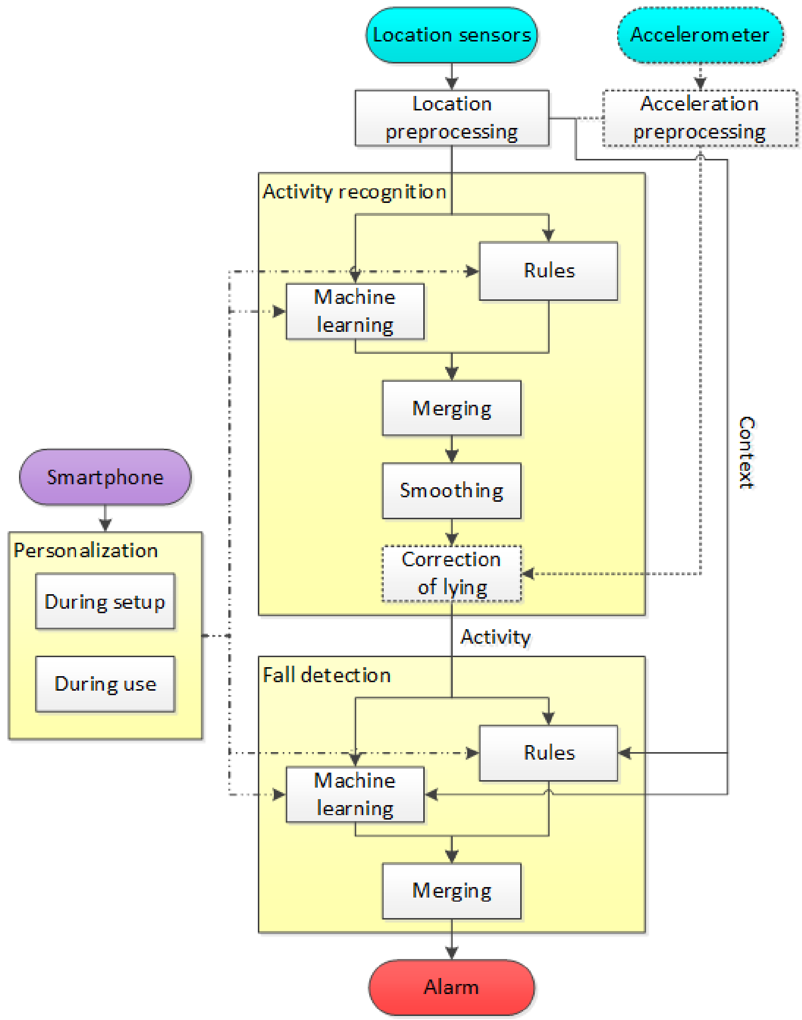

Figure 1: The Confidence fall-detection pipeline.

Sensors. Any real-time location system accurate enough to distinguish between different postures and activities can be used by the Confidence system. We used the Ubisense system, which employs ultra-wideband technology to determine the locations of up to four radio tags worn on the user's chest, waist and both ankles with sensors installed in the apartment. The number of tags can be reduced, but the chest tag is mandatory, since this location proved best for fall detection[8]. If not all four tags are used, the locations of the missing tags are estimated. The locations are sampled at 10 Hz. A three-axial accelerometer worn on the chest can be used optionally, also sampled at 10 Hz. We used Xsens accelerometers with the range of 5 g in all the experiments in the article, but any moderately accurate accelerometer would be acceptable.

Preprocessing. Location data needs to be preprocessed[11] because it is often very noisy. First, a median filter is applied, which eliminates large short-term changes in tag locations due to noise. Second, a filter that enforces anatomic constraints is used. It corrects errors such as an apparent lengthening of a limb. Third, the Kalman filter is applied, which smoothes sharp changes in both location and speed. Acceleration data is only low-pass filtered.

Activity Recognition. Filtered tag locations are used for activity recognition[11], which is performed by a machine-learning module and a rules module. The system recognizes basic activities: walking/standing, sitting, sitting on the ground, lying, the process of sitting/lying down, the process of standing up, and falling. The machine-learning module first represents the sensor data with attributes such as the tag velocities and the distances between the tags. These are fed into a classifier trained with the Random Forest machine-learning algorithm, which outputs the user's activity (with the accuracy of 74.1 %). The rules module employs similar attributes, except that expert-crafted rules are used to determine the user's activity.

The final activity is selected as the most probable given the outputs of the two modules, as computed from the true activities and the modules' outputs on labeled recordings. It is smoothed with a Hidden Markov Model, which eliminates infeasible activity transitions, for example from lying to standing without standing up in between. If an accelerometer is used, the system checks whether the chest tag is low and the orientation of the accelerometer close to horizontal. If so and the recognized activity is not lying, it is changed to lying. This correction is made because the recognition of lying is the most critical for fall detection.

Fall Detection. The user's activity and the context are finally combined to detect falls[12]. The context consists of the user's location in the apartment (on the bed, on the chair), which is provided by the location sensors, and whether the user is moving, which is provided by either the location sensors or the accelerometer. We consider an event a fall if the user falls and does not get up for 10 seconds. This should not be confused with the falling activity, which is not detected very reliably and is thus only one of the indicators of a fall.

Like the activity recognition, the fall detection is performed by a machine-learning and a rules module. The machine-learning module uses as attributes whether the user is on the bed or chair, whether the user is moving, how long ago the last falling activity was detected, and the percentages of the user's activities in the last 10 seconds. The latter are included because the activities preceding and following a potential fall can confirm or disconfirm it. These attributes are fed into two classifiers trained with the SVM and C4.5 machine-learning algorithms to classify the situation as a fall or non-fall. Since both classifiers are prone to false alarms, the module declares that a fall has occurred only if both classifiers output fall. The rules module contains several expert-crafted rules that check for the falling activity, whether the user is lying or sitting outside the bed or chair, and whether the user is moving.

We declare that a fall has occurred if either both the machine-learning and the rules module output fall, or if one of them outputs fall continuously for 3 seconds.

Personalization. Since the Confidence system is intended for various users whose behavior and environment can change over time, it is personalized during setup and use. During setup (or afterwards if needed), a technician indicates areas safe for lying, and makes a short recording of the user standing, sitting and lying. During use, the activity recognition is continuously adapted by means of semi-supervised learning[13], and the fall detection is adapted by taking into account manually triggered alarms and alarms for which the user indicated they were false[12] – a simple smartphone application with two buttons is used for this purpose.

3.2 Accelerometer-Based Fall Detection

To detect falls with accelerometers, we used two sensor configurations: (1) one accelerometer on the chest, which is fairly typical for accelerometer-based fall detection and is comparable with the Confidence system with one tag; and (2) three accelerometers on the chest, thigh and ankle, which is comparable with the Confidence system with four tags (the fourth accelerometer proved unnecessary[8]). The accelerations were sampled at 10 Hz, which proved sufficient and is also the frequency used by the location sensors.

The acceleration pattern during a fast fall is a decrease in the acceleration followed by an increase. This is because an accelerometer at rest registers the Earth's gravity and thus reads 1 g, while an accelerometer in free fall reads 0 g. When a person is moving normally, the acceleration is thus around 1 g. When a person starts falling, the acceleration decreases to around 0.5 g (perfect free fall is never achieved). Upon the impact with the ground, a short strong increase in the acceleration is measured, after which the acceleration returns to 1 g.

Threshold and Posture Rules. To detect falls with a threshold, we used the length of the acceleration vector from the chest accelerometer. The minimum and the maximum acceleration within a 1-second window were measured. If the difference between the maximum and the minimum exceeded the threshold of 1 g and the maximum came after the minimum, this was considered an indicator of a fall.

We improved the threshold-based fall detection with the measurement of the person's posture after a potential fall. Since accelerometers register the Earth's gravity, they can provide the orientation with respect to the ground. With the chest accelerometer only, the rules were designed so that a fall was detected when the acceleration exceeded the threshold and the chest was not upright. Three accelerometers allow a better recognition of the posture. In this case, a fall was detected when the acceleration exceeded the threshold and the person was either (1) lying or (2) sitting on the ground.

Threshold and Posture Rules with the Context. Finally, we took the context in terms of the location in the apartment into account. The location was measured with a single location sensor. With the chest accelerometer only, a fall was detected when the person was not on the bed or chair, and either (1) the threshold was exceeded or (2) the chest was not upright. With three accelerometers, a fall was detected when the person was not on the bed or chair, and either (1) the threshold was exceeded, (2) the person was lying or (3) the person was sitting on the ground. The main difference from the rules without the context is that exceeding the threshold is no longer mandatory for a fall to be detected, since lying or sitting outside the bed or chair can be detected. This allows the detection of slow falls.

Machine Learning. To detect a fall pattern with machine learning, a sliding window was used to transform the stream of acceleration data into instances for machine learning. The data could come from either one or three accelerometers. The instances were composed of a number of numerical attributes describing the data within each 0.8-second window , such as the average, minimum and maximum acceleration, and the accelerometer orientations[7]. They were fed into a classifier trained with the Random Forest machine-learning algorithm, which output fall or non-fall.

Machine Learning with the Context. Fall detection with machine learning was also augmented with the context in terms of the location in the apartment. In this case we declared that a fall had occurred when the classifier detected a fall and the person was not on the bed or chair.

4 Experiments

We trained the classifiers in the Confidence system and the classifiers for the accelerometer-based fall detection on a dataset composed of 5 hours of recordings containing the activities recognized by the Confidence system and three types of falls (available online: http://archive.ics.uci.edu/ml/datasets/Localization+Data+for+Person+Activity). This dataset was also used to select the best algorithms for machine learning (implemented in the Weka suite[14]) and to tune all the parameters of all the fall-detection methods.

We compared the performance of the fall detection methods in a laboratory on a scenario recorded both by 10 healthy young volunteers who did not participate in the training recordings (aged 24–31, 4 female and 6 male) and by 10 elderly volunteers (aged 66–80, 7 female and 3 male), 5 times by each, for a total of 24 hours of recordings. A medical expert advised how the young volunteers should behave to appear to be elderly. The recordings were made using the Ubisense real-time location system and Xsens accelerometers simultaneously. We also performed usability and acceptability tests, and long-term validation of the Confidence system in a real-world setting with elderly users. Ethical approval was obtained and the participants in the experiments gave informed consent.

4.1 Laboratory Test Scenario

We designed the scenario specifically to investigate events that may be difficult to recognize as falls or non-falls. It contains nine events listed in Table 1, examples of which can be viewed here: http://dis.ijs.si/confidence/ieee_percomp.html. We report the results by the elderly for all but two events: tripping and lying down quickly. We only report the results by the young for these two events because the elderly could not perform them safely in a realistic fashion.

No.	Description	Fall	Elderly
1	Sitting down normally on the chair	No	Yes
2	Tripping, landing flat on the ground	Yes	No
3	Lying down normally on the bed	No	Yes
4	Falling slowly (trying to hold onto furniture), landing flat on the ground	Yes	Yes
5	Sitting down quickly on the chair	No	Yes
6	Falling when trying to stand up, landing sitting of the ground	Yes	Yes
7	Lying down quickly on the bed	No	No
8	Falling slowly when trying to stand up (trying to hold onto furniture), landing sitting of the ground	Yes	Yes
9	Searching for something on the ground on all fours any lying	No	Yes

Table 1: Fall and non-fall events in the test scenario.

We selected representative fall types in consultation with medical experts. Since the related work demonstrated that accelerometers can detect typical fast falls accurately, we included only one such fall (event number 2). We included three atypical falls (4, 6 and 8) to test the use of contextual information, namely that a person is not expected to lie or sit on the ground (as opposed to the bed or chair). They are atypical in speed (4 and 8) and starting/ending posture (6 and 8). We included two events (5 and 7) that involve high acceleration and could thus be misclassified as falls by accelerometers. We also included an event (9) that involves voluntary lying on the ground, which could mislead the methods that rely on the information on the posture and location. The final two events (1 and 3) are perfectly normal and were included to verify that all the methods work correctly and do not recognize them as falls. Events similar to all of those in Table 1 were also present in the training recordings, except for the falls 6 and 8.

4.2 Laboratory Results

We compare the performance of the Confidence system and the accelerometer-based methods in Table 2. The left side of the table shows the results of the Confidence system (described in Section 3.1) with one tag and with the full complement of four tags, without and with the optional accelerometer. The right side shows the results of the accelerometer-based methods (described in Section 3.2) with one and three accelerometers. We present the results of the threshold and posture rules and machine learning (ML), without and with the location context (each method under a separate heading in Section 3.2). To compare the methods using purely location sensors vs. accelerometers, one should compare Confidence 1 or 4 tags w/o acc. vs. 1 or 3 accelerometers w/o context. To observe the benefit of adding an accelerometer to the location-based fall detection, one should compare Confidence 1 or 4 tags w/o acc. vs. Confidence 1 or 4 tags w/ acc. To observe the benefit of adding the location context to the accelerometer-based fall detection, one should compare 1 or 3 accelerometers w/o context vs. 1 or 3 accelerometers w/ context. To compare the methods using both sensor types but prioritizing location sensors vs. accelerometers, one should compare Confidence 1 or 4 tags w/ acc. vs. 1 or 3 accelerometers w/ context.

	Confidence				Accelerometer-based							
	1 tag		4 tags		1 accelerometer				3 accelerometers			
	W/o acc.	W/ acc.	W/o acc.	W/ acc.	W/o context		W/ context		W/o context		W/ context	
					Rules	ML	Rules	ML	Rules	ML	Rules	ML
Falls												
Tripping (2)	93.9	100.0	100.0	100.0	100.0	87.8	100.0	87.8	100.0	89.7	100.0	89.7
Falling slowly (4)	93.3	100.0	83.3	100.0	20.0	54.4	93.3	54.4	20.0	59.9	93.3	59.9
Falling sitting (6, 8)	98.3	100.0	98.3	100.0	1.7	61.7	31.7	61.7	8.4	66.0	96.7	66.0
Non-falls												
Lying quickly (7)	100.0	100.0	100.0	100.0	34.0	8.2	100.0	100.0	34.0	8.3	100.0	100.0
Sitting quickly (5)	93.3	88.3	93.3	95.0	100.0	2.3	96.7	100.0	100.0	5.6	96.7	100.0
On the ground (9)	86.7	90.0	96.7	97.7	100.0	35.6	53.3	35.6	100.0	36.9	96.7	36.9
Average	94.2	96.4	95.3	98.6	59.3	41.7	79.2	73.3	60.3	44.5	97.2	75.4
Other false alarms	1	1	4	1	2	3	2	3	1	1	1	1

Table 2: Fall detection accuracy (in %) of the Confidence system and of the accelerometer-based methods; accuracies 80–100% are shaded in green, 50–80% in yellow and 0-50% in red; event numbers from Table 1 are in parentheses.

Taking a look at the performance of the Confidence system in Table 2, we can see that all the accuracies are quite high, and the different versions of the system are close to each other. Four tags slightly outperformed one tag, except in the case of falling slowly, where one tag had the advantage of not being confused by the person's unusual posture after some of the falls. The version with the accelerometer slightly outperformed the version without it.

The performance of the accelerometer-based methods strongly depended on both the method and the event being recognized. The first event in Table 2, tripping, is a typical fast fall that was – as expected – recognized fairly accurately by all the methods. The machine learning made more mistakes than the threshold and posture rules.

Falling slowly was recognized poorly with the threshold and posture rules because the falling was too slow; only the context helped, because it allowed recognizing that the person was lying outside the bed. The machine learning also had difficulties with the slowness of the fall, although it still detected the falling pattern in most cases. It could not take advantage of the context, though, because if the falling pattern was not recognized, the information that the person was not on the bed or chair was useless. The results for falling sitting are similar, except that with one accelerometer the rules could not recognize it even with the context. The reason is that one accelerometer was not enough to distinguish sitting on the ground from other postures in which the chest was upright.

Lying and sitting quickly were easy to recognize for all the methods using the context, since they were taking place on the bed or chair. When the context was not used, all methods with one exception were misled by the high acceleration. The exception were the threshold and posture rules with one accelerometer, which correctly recognizing sitting quickly, because the upright posture indicated that no fall had occurred.

Searching on the ground was recognized well with the threshold and posture rules using one accelerometer, because there was no large acceleration, and with the rules using three accelerometers, because they could recognize that the person was on all fours and not lying. The rules using one accelerometer and the context were misled by the similarity with lying outside the bed, while the machine learning performed unexpectedly poorly.

4.3 Real-life Experience

Even though the laboratory tests of the Confidence system showed accurate fall detection, this is not enough to conclude the system is suitable for real-life use. Therefore we investigated the usability and acceptance by real older users, and conducted a long-term validation in two sheltered-housing facilities.

The usability and acceptance were tested by 12 participants (aged 65–77, 7 female and 5 male), who performed representative activities of daily living and interacted with the system through the smartphone user interface for around two hours each. The physical burden caused by wearing the tags was evaluated with the Borg's Rating of Perceived Exertion Scale and the similar Borg's CR10[15]. Most participants perceived no exertion when wearing the tags. The ease of use of the user interface was assessed through task-performance time and accuracy. After some adjustments based on the users' feedback, the interface consisted only of a large red button to raise the alarm and a smaller green button to cancel it. Task performance with such a simple interface was reasonably fast and fluid with no procedural errors. However, some users disliked the touch screen. The acceptance of the system was investigated by a semi-structured interview. Overall, participants considered that they would accept using the system. They agreed with the positive motives for using the system, such as making them feel safer and more independent living at home. They mostly disagreed with the potential barriers to the use of the system, such as the fear that it would be unreliable, or that they would feel stigmatized by it. They pointed out that the price could be a barrier: they estimated it at 1,000 EUR, which they would not be willing to pay, but they thought the public healthcare should – not an unexpected attitude considering the interview took place in Finland. They also indicated that the ergonomics and appearance would have to be improved, which is to be expected since the Confidence system was still a research prototype.

The Confidence system underwent long-term validation in two sheltered-housing facilities by 10 participants (aged 70–93, 8 female and 2 male). In one of the facilities the users lived very independently, whereas in the other they received more attention from the staff. The users wore all four tags for roughly eight hours per day for one month each. The system covered the room in their homes where they spent most of this time. No falls occurred during the test, so no true alarms could be raised. False alarms occurred on 20% of the recorded days. This is better than the results of a real-life experiment reported in the literature[2], in which multiple false alarms were raised during a 24-hour period. Like in the usability and acceptability test, the users felt that the ergonomics and appearance of the Confidence system should be improved. The tags were somewhat difficult to attach, but only one user complained that they were uncomfortable to wear. The smartphone interface posed no problems. The general attitude at the facility with the more independent users was positive. The users at the other facility felt the benefits did not outweigh the inconvenience, since they were frequently attended by the staff.

5 Conclusion

Falling is an important problem in elderly care, which has been often tackled using accelerometers. Most experimental results reported in the literature suggest that accelerometers are quite adequate for the task. However, we suspect that most of the experiments involved typical falls with high acceleration upon the impact with the ground, and contrasted them with typical daily activities. Slow falls are more difficult to detect, while atypical fall-like events can trigger false alarms that limit the acceptability of fall-detection solutions. We experimented with just such events, showing that accelerometers alone indeed cannot distinguish falls from non-falls successfully, so improvements are needed.

Our experiments with the Confidence system and accelerometer-based fall detection led us to the following key conclusions:

- The location context helps the most: the Confidence system and the accelerometer-based methods using the context strongly outperformed the methods without it. The test scenario was admittedly designed to investigate events whose recognition was expected to be difficult without the context, but that does not mean that such events are rare in real life.
- After the context, accurate recognition of the posture is key: of the methods using the context, the Confidence system and the threshold and posture rules using three accelerometers performed best. What is common to these methods is accurate recognition of the posture.
- A combination of accelerometers and location sensors is better than either sensor type alone: the top two performers were the Confidence system with four tags and an accelerometer, and the threshold and posture rules using three accelerometers with the context provided by a location tag.
- When using accelerometers, the threshold and posture rules outperformed machine learning. Our hypothesis had been that machine learning would capture the subtleties of falling patterns to achieve results comparable to the rules. To test it, we used attributes for machine learning describing accelerations during falling, instead of attributes such as those in the rules. The hypothesis was not confirmed, possibly because we did not have enough training data to capture the large variation in falls. The rules did not suffer from this problem because they reflected expert understanding of falling based on extensive experience[13].
- The Confidence system appears to be sufficiently accurate in real life, usable and acceptable by the elderly who live independently. Accelerometers are physically similar to the Confidence tags, so they should also be acceptable; the only concern is their higher energy consumption. The best balance between the accuracy and comfort is probably achieved by the Confidence system with one tag and one accelerometer in a single enclosure (on the chest). In this case the average

fall-detection accuracy is 96.4%, which is only 2.2 percentage points worse than the most accurate solution that uses four tags.

While our experiments were performed on a limited number of event types and people, we believe the conclusions apply more widely. It is reasonable to expect that accelerometers – particularly if focusing on the fall impact – will have difficulty with most slow falls (such as falling slowly and sitting – 4, 6, 8), especially if the person ends up in a partially upright posture (falling sitting – 6, 8). They will also have difficulty with events where the motion is fall-like, but the context suggests it is safe (lying quickly – 7). It is also reasonable to expect that knowing whether a potential fall occurred on the ground vs. the bed or chair should resolve most of these difficulties. It may, however, lead to difficulties with safe activities on the ground (searching on the ground – 9). This could be a problem for some users, although those who are at the greatest risk from falling due to frailty typically avoid getting on the ground because they find it difficult to get up again. An issue not addressed in this article are falls that occur on the bed, chair or other location that is considered safe.

Considering the importance of the context, we will explore additional contextual information in the future. For example, the user's high-level activity would be useful because fall-like events are much more common while exercising than while watching TV. Another interesting avenue of research is a tighter integration of the location sensors and accelerometers, since the combined approach proved best in our experiments. Finally, different approaches to location sensing will have to be considered for a practical application, since the cost of the Ubisense system is currently too high for widespread adoption. However, the strong commercial interest in indoor positioning will probably soon lead to more accessible technologies suitable for our purposes.

Acknowledgement

The research leading to these results has received funding from the European Community's Framework Programme FP7/2007–2013 under grant agreement nº 214986. Operation part was also financed by the European Union, European Social Fund.

References

1. D. Wild, U.S. Nayak, B. Isaacs, "How Dangerous are Falls in Old People at Home?," *British Medical J. (Clinical Research Ed.)*, vol. 282, no. 6260, 1982, pp. 266–268.
2. F. Bagalà, C. Becker, A. Capello, L. Chiari, K. Aminian, J.M. Hausdorff, W. Zijlstra, and J. Klenk, "Evaluation of Accelerometer-Based Fall Detection Algorithms on Real-World Falls," *PLoS One*, vol. 7, no. 5, 2012.
3. S. Shan, and T. Yuan, "A Wearable Pre-Impact Fall Detector Using Feature Selection and Support Vector Machine," *Proc. IEEE 10th Int'l Conf. on Signal Processing* (ICSP 10), 2010, pp. 1686–1689.

4. M. Yuwono, S.W. Su, and B. Moulton, “Fall Detection Using a Gaussian Distribution of Clustered Knowledge, Augmented Radial Basis Neural-network, and Multilayer Perceptron,” *Proc. 6th Int’l Conf. Broadband Communications & Biomedical Applications* (IB2Com 11), 2011, pp. 145–150.
5. Q. Li, J.A. Stankovic, M.A. Hanson, A.T. Barth, J. Lach, and G. Zhou, G., “Accurate, Fast Fall Detection Using Gyroscopes and Accelerometer-Derived Posture Information,” *Proc. 6th Int’l Workshop Wearable and Implantable Body Sensor Networks* (BSN 09), 2009, pp. 138–143.
6. Q. Li, and J.A. Stanković, “Grammar-Based, Posture- and Context-Cognitive Detection for Falls with Different Activity Levels,” *Proc. 2nd Conf. Wireless Health* (WH 11), 2011, article 6.
7. M. Luštrek, H. Gjoreski, S. Kozina, B. Cvetković, V. Mirchevska, and M. Gams, “Detecting Falls with Location Sensors and Accelerometers,” *Proc. 23rd Conf. Innovative Applications of Artificial Intelligence* (IAAI 11), 2011, pp. 1662–1667.
8. H. Gjoreski, M. Gams, and M. Luštrek, “Context-Based Fall Detection and Activity Recognition Using Inertial and Location Sensors,” *J. Ambient Intelligence and Smart Environments*, vol. 6, no. 4, 2014.
9. A. Kashevnik, and M. Shchekotov, “Comparative Analysis of Indoor Positioning Systems Based on Communications Supported by Smartphones,” *Proc. 12th Conf. FRUCT Assoc.*, 2012, pp. 43–48.
10. M.E. Bowen, J. Craighead, C.A. Wingrave, and W.D. Kearns, “Real-Time Locating Systems (RTLS) to Improve Fall Detection,” *Gerontechnology*, vol. 9, no. 4, 2010, pp. 464–471.
11. B. Kaluža, B. Cvetković, E. Dovgan, H. Gjoreski, M. Gams, V. Mirchevska, and M. Luštrek, “Multiagent Care System to Support Independent Living,” *Int’l J. Artificial Intelligence Tools*, vol. 23, no. 1, 2014.
12. V. Mirchevska, M. Luštrek, and M. Gams, “Combining Domain Knowledge and Machine Learning for Robust Fall Detection,” *Expert Systems*, 2013, preprint, 6 Mar. 2013, doi: 10.1111/exsy.12019.
13. B. Cvetković, B. Kaluža, M. Luštrek, and M. Gams, “Adapting Activity Recognition to a Person with Multi-Classifier Adaptive Training,” to be published in *J. Ambient Intelligence and Smart Environments*, 2014.
14. M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. H. Witten, “The WEKA Data Mining Software: An Update,” *SIGKDD Explorations*, vol. 11, no. 1, 2009, pp. 10–18.
15. G. Borg, *Borg's Perceived Exertion and Pain Scales*, Human Kinetics, 1998.

Author Bios

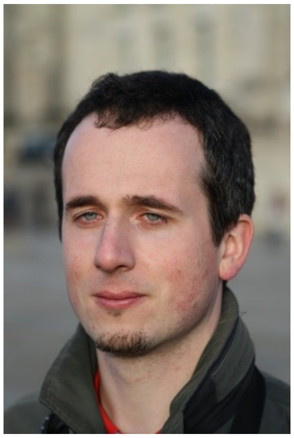

Mitja Luštrek is the head of the Ambient Intelligence Group at the Department of Intelligent Systems at Jožef Stefan Institute. His main research interest is ambient intelligence, particularly the analysis of human behavior using sensor data. He holds a Ph.D. degree in computer and information science from the University of Ljubljana. Email: mitja.lustrek@ijs.si, postal address: Jožef Stefan Institute, Jamova cesta 39, 1000 Ljubljana, Slovenia.

Hristijan Gjoreski is a researcher at the Department of Intelligent Systems at Jožef Stefan Institute. His research interests include context-based reasoning, wearable computing and ambient intelligence. He holds a M.Sc. degree in computer science from the Jožef Stefan International Postgraduate School. Email: hristijan.gjoreski@ijs.si, postal address: Jožef Stefan Institute, Jamova cesta 39, 1000 Ljubljana, Slovenia.

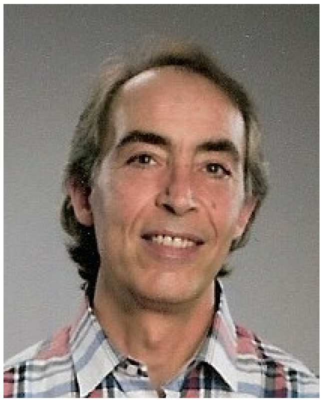

Narciso González Vega has been co-leader of the Service Science Laboratory at Agora Center at the University of Jyväskylä. His main research interest rests on the interface between older people and technologies which can expand healthy life years. He holds a Ph.D. degree in Psychology from the University of Jyväskylä. Email: narciso.gonzalez-vega@jyu.fi, postal address: Agora Human Technology Center, University of Jyväskylä, Mattilanniemi 2, 40014 Jyväskylä, Finland.

Simon Kozina is a researcher at the Department of Intelligent Systems at Jožef Stefan Institute. His main research interest is ambient intelligence. He holds a B.Sc. degree in computer and information science from the University of Ljubljana. Email: simon.kozina@ijs.si, postal address: Jožef Stefan Institute, Jamova cesta 39, 1000 Ljubljana, Slovenia.

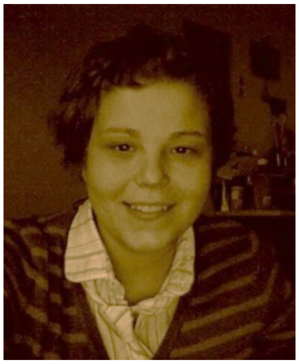

Božidara Cvetković is a researcher at the Department of Intelligent Systems at Jožef Stefan Institute. Her main research interests are semi-supervised machine learning and personalization in the area of ambient intelligence. She holds a B.Sc. degree in computer and information science from the University of Ljubljana. Email: boza.cvetkovic@ijs.si, postal address: Jožef Stefan Institute, Jamova cesta 39, 1000 Ljubljana, Slovenia.

Violeta Mirchevska is a researcher at the Department of Intelligent Systems at Jožef Stefan Institute. Her main research interest is ambient intelligence. She holds a Ph.D. degree in computer and information science from the Jožef Stefan International Postgraduate School. Email: violeta.mircevska@ijs.si, postal address: Jožef Stefan Institute, Jamova cesta 39, 1000 Ljubljana, Slovenia.

Matjaž Gams is the head of the Department of Intelligent systems at Jožef Stefan Institute, and professor at the University of Ljubljana and Jožef Stefan Postgraduate School. His research interests include ambient intelligence, machine learning, agents, hybrid learning and reasoning. He holds a Ph.D. degree in computer and information science from the University of Ljubljana. Email: matjaz.gams@ijs.si, postal address: Jožef Stefan Institute, Jamova cesta 39, 1000 Ljubljana, Slovenia.