Automatic In-Door Fall Detection based on Microwave Radar Measurements

Peter Karsmakers*§, Tom Croonenborghs†§, Marco Mercuri‡, Dominique Schreurs‡ and Paul Leroux‡¶
*KU Leuven, Div. ESAT-SISTA, Kasteelpark Arenberg 10, B-3001 Heverlee, Belgium
†KU Leuven, Div. CS, Celestijnenlaan 200A, B-3001 Heverlee, Belgium
‡KU Leuven, Div. ESAT-TELEMIC, Kasteelpark Arenberg 10, B-3001 Heverlee, Belgium
¶KU Leuven, Div. ESAT-MICAS, Kasteelpark Arenberg 10, B-3001 Heverlee, Belgium
§KH Kempen University College, Div. IBW, Kleinhoefstraat 4, B-2440 Geel, Belgium
{Peter.Karsmakers,Marco.Mercuri,Dominique.Schreurs,Paul.Leroux}@esat.kuleuven.be
Tom.Croonenborghs@cs.kuleuven.be

***Abstract*—The use of a Continuous Wave (CW) Doppler radar is proposed for non-invasive automatic detection of human falls. This radar technology can be used since fall incidents can be characterized by changes in speed. In this paper we show that speed measurements obtained from different activities, using a radar fixed on the ceiling, can automatically discriminate between fall incidents and other activities with good accuracy. The activities we consider are falling, walking, running, and sitting. Off-the-shelf machine learning techniques are used to estimate an activity classification model.**

***Index Terms*—Fall detection, activity classification, health monitoring, radar, remote sensing.**

I. Introduction

Statistics show that 30% to 45% of the persons aged 65 or older living at home and more than 50% of the elderly living in a nursing home fall at least once a year [1]. These fall incidents cause severe injuries in 10% to 15% of the cases. The lack of timely aid can even lead to more severe complications (e.g. dehydration, pressure ulcers and even death). Although not all falls lead to physical injuries, psychological consequences are equally important, leading to fear of falling, losing self-confidence and fear of losing independence [1]. Taking the ongoing aging of the population into account, it is evident that the automatic detection of fall incidents is getting more and more important.

The existing detectors are mostly based on wearable sensors. However a market study of SeniorWatch [2] showed that the sensors are not worn at all times (e.g. at night). Furthermore, when the device needs to be operated by a button, such as e.g. a Personal Alarm System, the person is often unable to activate the alarm system due to the resulting confusion of a fall incident and the complexity of making the alarm call. A remote monitoring approach based on radar principles can overcome these disadvantages. Current systems under investigation for contactless fall detection are mostly based on video cameras [3]. In contrast to these systems, the radar based approach might alleviate some privacy concerns.

Radars are used before for remote health monitoring, e.g. a CW Doppler radar [4] or UWB IR radar in [5]. These papers however focus on contactless vital signs detection, especially heartbeat and respiration rate. Similar work is found in [6] were falls are detected using a different radar setup and classification strategy as proposed here.

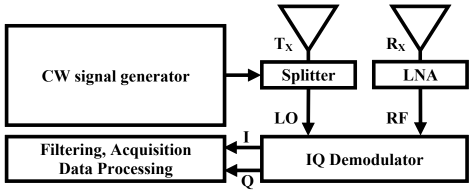

Figure 1: Radar architecture.

II. Radar Architecture

In this section we introduce the employed radar architecture. Since a fall can be considered as a movement with varying speed, a CW radar exploiting the Doppler effect is adopted. The block diagram is presented in Fig. 1. We use a Quadrature CW Doppler radar as having both I and Q signals improves the accuracy of the measurement. The I and Q samples are related as $C = I + jQ$. The architecture consists of two antennas, a power slitter, an LNA, and the IQ demodulator. The RF frequency is 1.7 GHz and the power sent to the transmit antenna is -10dBm.

A set of samples C results in a signal whose frequency is altered in proportion to the velocity of the target according to the Doppler effect: $f_d(t) = \frac{2fv(t)}{c}$ where f is the frequency of the incident signal (1.7GHz), $v(t)$ is the velocity of the moving object, c is the speed of light, and $f_d(t)$ is the resulting shift in frequency.

In order to demonstrate the functionality of the classification algorithm, the radar is realized as a board design using off-the-shelf components (see Fig. 2). A more compact design adopting a single antenna is under development. Fig. 2 shows the experimental setup with the antennas mounted on the ceiling. Note that the metal shelves do not influence the signal. In fact, the received signals are AC coupled and just movements can be detected. Fig. 3 and Fig. 4 show the measured I and Q responses during a fall and a walk respectively. The time variation of the I and Q signals is proportional to the velocity of the target. In the first case of the fall incident, Fig. 3 clearly shows how the frequency signal (and thus the speed) increases

Figure 2: Radar setup. The antennas are fixed on the ceiling while the electronics are positioned on the shelf. The I and Q signals are acquired by an oscilloscope. The electronic boards were marked.

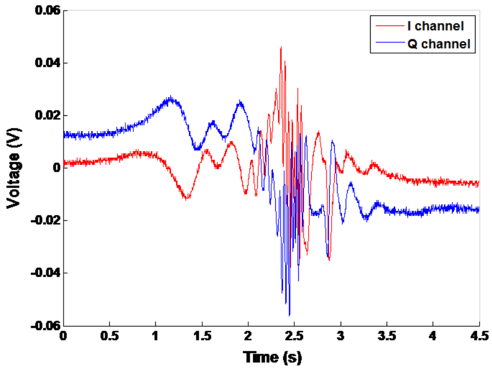

Figure 3: Speed signal during a fall. The frequency of the signal is proportional to the velocity of the person during the fall.

with time and then abruptly stops when the person hits the ground. In the second case of the walking activity (Fig. 4), the frequency of the signal is proportional to the speed and position of the person. In the middle of Fig. 4 the frequency is low since the person is passing underneath the radar which is mounted on the wall. In our lab experiments, we used an inflatable mattress when performing the fall activity. This can be seen in Fig. 3, since the signal does not stop suddenly but there is the effect of the rebounds on the mattress. These signals are used to detect a fall and to distinguish it from a normal movement (i.e., walking, sitting, and running) as will be explained in the next section.

III. Automatic Fall Detector

In this paper we propose an approach to automatically classify radar-based data segments into activities. The primary goal is to discriminate falls from other movements. However, to show that a more detailed discrimination is possible, we will employ a model that classifies the radar data into four different activities: falling, walking, running, and sitting.

Such a model can for example be constructed by making a list of rules guided by human intuition (e.g. during a fall, the

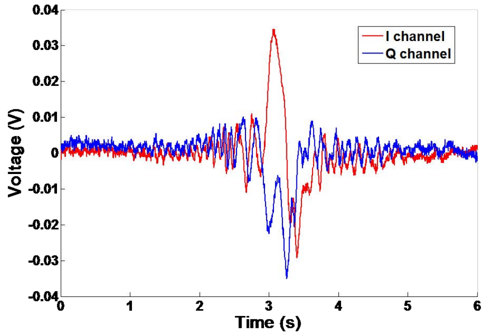

Figure 4: Speed signal during a walk. The frequency of the signal is proportional to the velocity and position of the person during the walk.

speed continuously increases until the floor is hit). However, compiling a set of rules can be difficult and time consuming. We therefore opt to explore the use of an off-the-shelf machine learning technique which automatically discovers structure (rules) by inspecting example data. A common approach is to learn this structure in a supervised manner meaning that annotated examples need to be presented to the learning algorithm [7].

Note that in this work we assume that segments, containing radar measurements of a single activity, are available. Making the system work in this situation is considered as a first test before making the extension to recognize activities from a continuous stream of radar data.

A. *Supervised activity classification*

The machine learning community exhibits a large set of methods for supervised learning. We have chosen to use a specific framework out of the family of kernel methods named Least Squares Support Vector Machines (LS-SVM) [8]. This framework has shown to give good results on a wide range of applications ([9]) and has the advantage that it can be altered easily to be used on different kinds of data (static reals, timeseries, ...).

To illustrate this activity classification, consider a binary classification task where the goal is to find a suitable decision line between two activities. With suitable it is meant that the learning algorithm needs to obtain models with adequate generalization abilities so that it performs well on predicting class labels for unseen examples, i.e. new measurements. To learn a classification model, a set of N observations $\mathbb{D} = \{(x_i, y_i)\}_{i=1}^{N}$ of measurements $x_i \in \mathbb{R}^D$ and the corresponding observed output values $y_i \in \{-1, 1\}$ is used. When using $D = 2$, i.e. the measurement are represented as $x_i \in \mathbb{R}^2$, the learning problem can be visualized as shown in the left plot of Fig. 5[1]. Each class is represented by a specific marker. All the examples of the star class could be assigned 1 as a class label while the examples of the other class can be assigned the

[1]For instance, a radar segment can be represented by two features using the average frequency and energy of the signal.

label -1. As a second step a function $f(x)$ is estimated which maps the measurement x_i to the corresponding numeric labels y_i. Having such a function, one could predict the label of an unseen sample by $\hat{y}^* = \text{sign}(f(x^*))$ where $\text{sign}(\cdot)$ returns $+1$ for positive values of $f(\cdot)$ and -1 for negative values.

In the LS-SVM framework the mapping function is defined as $f(x) = w^T \varphi(x) + b$. This function is linear in the parameters $w \in \mathbb{R}^{\varphi_D}$ and $b \in \mathbb{R}$. These parameters are then optimized according to a convex optimization criterion meaning that the global optimum solution can always be found. Given a set of examples (called training set) the following convex primal objective can be formulated

$$\min_{w,e_i,b} \frac{1}{2} w^T w + \frac{\gamma}{2} \sum_{i=1}^{N} e_i^2, \tag{1}$$

where $e_i = y_i - f(x_i)$, and $\gamma \in \mathbb{R}^+ \setminus \{0\}$. In this objective, it is clear that minimizing the second term corresponds to making as few errors as possible. The first term controls the flexibility of the solution, e.g. large values in vector w may as a consequence give large shifts in $f(x)$ for small changes in x. The opposite is true for small values in w. This is related to the generalization ability of the solution, a too flexible solution might have a bad generalization performance (e.g. small changes due to measurement noise might give large unwanted changes in $f(x)$). Hence minimizing the first term in (1) requires a decrease in flexibility of the solution. Determining which term of the two is more important is problem dependent. Therefore a trade-off parameter (also called regularization parameter) γ is introduced. A large value for γ means that during the optimization process more attention is given to finding a parameter configuration that makes less errors at the cost of a more flexible, possibly less general, solution. On the other hand for small values of γ a less flexible solution is considered better at a cost of a higher error rate. Since this so called hyperparameter γ is application specific, it needs to be tuned for every application.

The above explained binary classification framework has many alternative multi-class extensions [8] where the classifier discriminates between multiple labels instead of only two different labels.

B. Kernel function

In the mapping function a feature map $\varphi(\cdot)$ is used. Such a feature map is visualized in Fig. 5. Data samples x in the input space are embedded into a vector space, called the feature space, as $\varphi(x)$. Then linear relations are sought, using well-known and stable methods, among the images of the data in the feature space.

Since finding and computing an appropriate $\varphi(\cdot)$ can be difficult, equation (1) is reformulated as a dual objective where the function $\varphi(\cdot)$ appears only in an inner product $\varphi(x)^T \varphi(x')$. This product can be replaced by a so called kernel function which computes inner products in the feature space directly from the inputs. Hence, the feature map is only implicitly defined. Any valid kernel function $K : \mathbb{R}^D \times \mathbb{R}^D \rightarrow \mathbb{R}$ corresponds with an inner product in a corresponding feature space as long as the function K is positive semi-definite. Popular choices are the linear kernel $K(x, x') = x^T x'$, and the Radial Basis Function (RBF) kernel $K(x, x') = \exp\left(-||x - x'||_2^2/\sigma^2\right)$ with kernel bandwidth σ strictly positive. This σ is another hyperparameter which needs to be tuned.

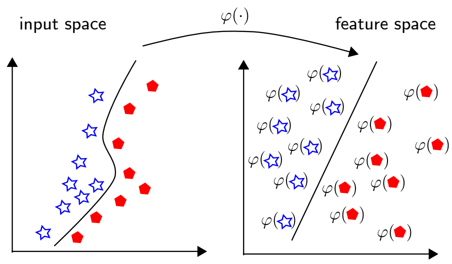

Figure 5: The function $\varphi(\cdot)$ embeds the data into a feature space where the non-linear decision line now appears linear. The kernel computes inner products in the feature space directly from the inputs.

The previous 'static' kernels are appropriate when each example (e.g. activity) can be represented by a vector. However, when the observations consist of sequences of vectors and activity labels $\{(\underline{X}_i, y_i)\}_{i=1}^N$ where $\underline{X}_i$ denotes a sequence of feature vectors, a different semi-positive definite kernel is needed which can handle sequences of vectors. One such example is the Global Alignment (GA) kernel developed in [10]. Using this kernel furthermore requires the tuning of a kernel bandwidth σ. The latter indicates the power of the considered learning framework. The classification model can be tailored easily to a specified application by choosing an appropriate kernel function.

IV. Experimental results

Using the setup described in Section II a data set was built containing 60 examples measured from 2 different persons each doing 4 predefined activities i.e. falling, walking, running and sitting at different locations in the recording room.

Before learning a model, the raw radar data is preprocessed. For each activity a radar segment of 2 seconds, meaning 768 samples since the downconverted waveform is sampled at a frequency of 384Hz, was selected. For this data, 2 seconds was found to be sufficient to cover the details of the activities.

Given such segments, the data is then transformed using a Fast Fourier Transform (FFT) from which only the magnitude spectrum is retained. We considered two alternatives: a) compute a 1024 FFT directly on the complete segment and b) use a Short Time FFT (STFT). In case of the latter the segment is first chopped into 50% overlapping frames which are each multiplied with a Hamming window after which a 64 point FFT is computed on each of these frames. Opposed to the FFT, the STFT can represent time dependent structures. This can be important since e.g. the activity sitting is characterized by a gradual increase in velocity followed by a gradual decrease. For the first option LS-SVM in combination with a linear and RBF kernel is used. When the STFT is used, per activity a sequence of vectors $\underline{X}_i$ exists. To be able to handle the latter, LS-SVM is combined with the GA kernel. An alternative method called Dynamic Time Warping (DTW)[11] combined with a Euclidean distance measure is frequently used to classify sequences of vectors. In order to compare the LS-SVM with GA solution to this standard method an additional experiment was included. Prior to the learning phase the data was standardized such that each dimension has zero mean and unit standard deviation.

In order to validate the classification models an independent test set (with data not used to learn the classification model) is needed. For this purpose the available data was split up into two parts. Each part contains the data of a single person. Next, a LS-SVM model was trained on a single partition and validated on the other. This process was repeated two times since the data of two test persons was available. In each run the hyperparameters were determined using a cross validation scheme [8]. The resulting classification accuracies are shown in Table I. It can be noticed that the GA kernel which incorporates time dependent information is outperforming the static kernels. This is as expected since the activities falling and sitting clearly exhibit a time dependent structure (e.g. in case of falling an increase of velocity is followed by a sudden stop). Moreover the GA kernel variant outperforms the DTW alternative. Due to the type of validation strategy we can conclude that for this experiment a classification model estimated on the data of one person is generalizable to be used on data of another person. The confusion matrix of the GA kernel in Table II presents the confusions between the different activities. Note that the predictions of the two LS-SVM models were combined. First, activities were predicted for the first person given a model estimated with data of the second person. Next, predictions for the second person computed using a model trained with data of the first person only were added. In the confusion matrix it is for instance seen that running is perfectly discriminated from the other activities and that a falling segment was only misclassified twice, once as walking and once as sitting.

Method	Accuracy
LS-SVM, Linear kernel	86.7%
LS-SVM, RBF kernel	85.0%
DTW	90%
LS-SVM, GA kernel	95.0%

Table I: Accuracy of correctly classified activity labels using different classification models.

V. Conclusion and Future Work

This paper shows the feasibility of using CW radar for remote automatic fall detection. On the acquired measurements an activity model is learned that can discriminate 4 activities (i.e. falling, walking, running, and sitting) with an accuracy of 95% in terms of correct classifications. The evaluation was carried out on measured data generated by a person of which the data was not used to learn the activity classification model. Future research will focus on testing this framework on a larger set of examples and persons.

	falling	walking	running	sitting
falling	13	1	0	1
walking	0	18	0	1
running	0	0	6	0
sitting	0	0	0	20

Table II: The confusionmatrix represents per cell the count of instances for which the row indicates the known activity and the column the predicted activity using the GA kernel.

VI. Acknowledgements

This work has been supported by FWO-Flanders and KU Leuven GOA project.

References

[1] K. Milisen, E. Detroch, K. Bellens, T. Braes, K. Dierickx, W. Smeulders, S. Teughels, E. Dejaeger, S. Boonen, and W. Pelemans, “Falls among community-dwelling elderly: a pilot study of prevalence, circumstances and consequences in flanders,” *Journal Gerontol Geriater*, vol. 35, no. 1, pp. 15–20, 2004.

[2] SeniorWatch, “Fall detector: Case study of european ist seniorwatch project,” *Tech. Rep. SeniorWatch*, 2001.

[3] G. Debard, P. Karsmakers, M. Deschodt, E. Vlaeyen, J. V. den Bergh, E. Dejaeger, K. Milisen, T. Goedemé, T. Tuytelaars, and B. Vanrumste, “Camera based fall detection using multiple features validated with real life video,” in *Workshop proceedings 7th international conference on intelligent environments - IE 2011*, vol. 10, 2011, pp. 441–450.

[4] A. Droitcour, O. Boric-Lubecke, V. Lubecke, J. Lin, and G. Kovac, “Range correlation and i/q performance benefits in single-chip silicon doppler radars for noncontact cardiopulmonary monitoring,” *IEEE Transactions on Microwave Theory and Techniques*, vol. 52, no. 3, pp. 838–848, 2004.

[5] I. Y. Immoreev, S. Samkov, and T.-H. Tao, “Short-distance ultra wideband radars,” *IEEE Aerospace and Electronic Systems Magazine*, vol. 20, no. 6, pp. 9–14, 2005.

[6] L. Liang, M. Popescu, M. Skubic, M. Rantz, T. Yardibi, and P. Cuddihy, “Automatic fall detection based on doppler radar motion signature,” in *In Proceedings of the 5th International Conference on Pervasive Computing Technologies for Healthcare (PervasiveHealth)*, Dublin, Ireland, 2011, pp. 222–225.

[7] C. M. Bishop, *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Secaucus, NJ, USA: Springer-Verlag New York, Inc., 2006.

[8] J. A. K. Suykens, T. Van Gestel, J. De Brabanter, B. De Moor, and J. Vandewalle, *Least Squares Support Vector Machines*. World Scientific, 2002.

[9] T. Van Gestel, J. Suykens, B. Baesens, S. Viaene, J. Vanthienen, G. Dedene, B. De Moor, and J. Vandewalle, “Benchmarking least squares support vector machine classifiers,” *Machine Learning*, vol. 54, no. 1, pp. 5–32, 2004.

[10] M. Cuturi, “Fast global alignment kernels,” in *In Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, 2011, pp. 929–936.

[11] L. Rabiner and B. Juang, *Fundamentals of speech recognition*. Prentice-Hall, 1993.