Combining machine learning and expert knowledge for classifying human posture

Violeta Mirčevska, Mitja Luštrek, Matjaž Gams

Jožef Stefan Institute, Department of Intelligent Systems
Jamova 39, SI-1000 Ljubljana, Slovenia
violeta.mircevska@ijs.si, mitja.lustrek@ijs.si, matjaz.gams@ijs.si

Abstract

This paper presents a rule engine for classifying human posture according to information about the location of body parts. The rule engine was developed by enriching decision trees with expert knowledge. Results show 5 percentage points improvement in accuracy compared to support vector machines and a significant 11 percentage points compared to decision trees. The incorporation of expert knowledge overcomes the problem of classifier over-fitting observed with classifiers induced with machine learning. Better robustness of the posture classification rule engine is expected in real-life tests in comparison to classifiers induced with machine learning.

1 Introduction

The goal of the European FP7 project CONFIDENCE – Ubiquitous Care System to Support Independent Living [1] is to develop a system that will monitor the health conditions of the elderly in real time [2] [3]. The reasoning of the system is based on the positions of the user's body parts. Positions are measured solely with the use of radio sensors. The key advantage of the system is that it guaranties no intrusion in the privacy of the user.

In order to monitor the health conditions of the elderly, the system needs to be able to make conclusions about the general state of the person and changes in his/her behavior and detect abnormal situations that may be caused by health problems. Accurate classification of the posture of the human is essential in order for these higher level conclusions to be valid.

Due to the wide variety of body configurations, it is difficult to record all possible situations and to obtain representative training dataset for posture classification. If the training dataset is not a representation of the observed problem domain, machine learning models may over-fit, because with machine learning techniques statistically relevant patterns in the training data are extracted. In order to overcome over-fitting, we considered enriching machine learning with expert knowledge. Since humans are good at imagining body structures and postures not represented in the training data, application of expert knowledge should improve the accuracy and robustness of the classifier.

In this paper we present a rule engine for classifying human posture which combines machine learning and expert knowledge. More precisely, we present its structure, the methodology according to which it was developed and the results achieved. Detailed description of posture classification with the use of machine learning techniques can be found in [4] [5] [6] [7].

The paper is organized as follows. Section 2 presents the architecture of the rule engine and the procedure by which rules were extracted. In Section 3 we present evaluation of the rule engine and compare it with other machine learning techniques. Section 4 concludes the paper with conclusions and open issues for future work.

2 Posture classification rule engine

The posture classification rule engine can recognize six postures - standing, sitting, lying, falling, moving downwards (normally) and moving upwards.

The classification is done based on the position of the neck and the ankles of the user. More precisely, only the Z-coordinate of the neck and ankles is considered.

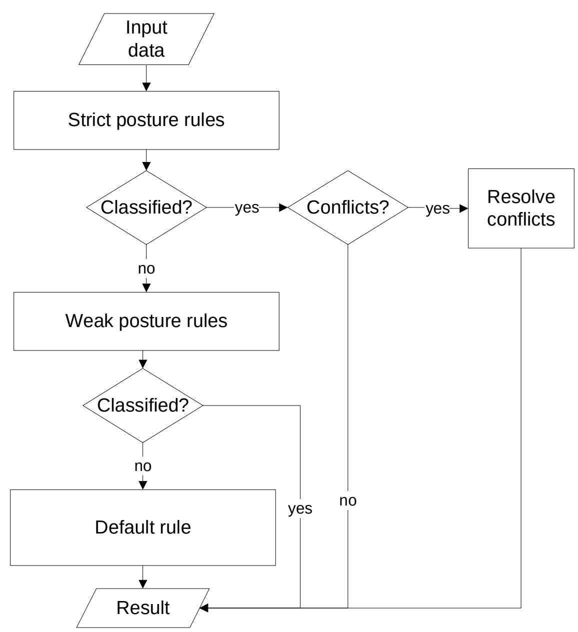

Figure 1: Architecture of the posture classification rule engine

The X- and Y- coordinates are not relevant, because they refer to the place in the room where the user is. Additionally, the neck-ankle distance in Z-direction and its projection on the XY-plane is used. These distances are most important for distinguishing between lying, sitting and standing. Finally, the velocity of the neck is considered. Being one of the topmost body parts, the velocity of the neck is highest during falls, moving downwards and upwards, making it suitable for distinguishing them.

The rule engine (Figure 1) is composed of three types of rules – strict posture rules, weak posture rules and default rule.

The strict posture rules contain precise definitions of the body configuration in each of the postures of interest. They have been created by enriching rules obtained from decision tree models with expert knowledge (Section 3). Each instance is first processed by the strict posture rules. If it is covered by set of strict posture rules describing only one posture class, this class in assigned to it. Conflicts when a particular instance is covered by rules of more than one posture are resolved as presented in Table 1. Conflicts appear between rules for adjacent classes, e.g. standing and going down. Since the rules for standing, sitting, lying and falling were constructed in a way that only pure postures are captured, they are chosen when there is a conflict with a rule for moving downwards/upwards.

Table 1: Resolution of conflicts among the posture rules

Conflict	Result
Standing and moving downwards/upwards	Standing
Sitting and moving downwards/upwards	Sitting
Lying and moving downwards/upwards	Lying
Falling and moving downwards/upwards	Falling

The weak posture rules specify the most probable class according to the user's neck-ankle distance. The weak posture rules were created by using expert knowledge. Each instance which is not covered by any of the strict posture rules is processed by the weak posture rules.

Finally, the default rule is used to assign a class to an instance that is not covered by both the strict and the weak posture rules. Since the current posture of a person is highly correlated with the posture he/she had in the previous time interval, the default rule assigns the class of the previous time interval to the instance in the current time interval.

2.1 Strict rule extraction procedure

The strict rules for classifying postures were extracted by a procedure similar to the one presented in [8] in the following way:

1. Create one-against-all dataset for a particular posture
2. Create decision trees
3. Extract rules with high precision and possibly high recall
4. Modify the extracted rules by expert knowledge.

One-against-all datasets were created for each posture. The idea is to concentrate on differences between a particular posture and all other postures. Then, decision trees were generated with the purpose of identifying the attributes that best separate the examples of this particular posture from all other postures.

The decision tree induction technique presents only the best hypothesis for the problem at hand because it performs general-to-specific hill-climbing search through the space of possible hypotheses. Because of this, relevant information may be hidden behind the best hypothesis. We constructed several decision trees using different attribute sets in order to overcome this problem. A decision tree was first built with all attributes. Then, the procedure was repeated by removing the attribute at the root node or attribute near the root node with the aim of finding relevant hidden hypotheses, until the classification accuracy of the resultant tree significantly dropped.

From the decision trees, rules with high precision and, possibly, high recall were extracted and modified by expert knowledge. The conditions of the extracted rules were made stricter, especially for the classes standing, sitting, lying and falling, improving their precision at the expense of recall. The aim was to correctly classifying pure postures, neglecting the borderline cases.

The rules were added to the set of rules for the particular class.

3 Experiments

The performance of the rule engine was tested and compared with four machine learning techniques – support vector machines, random forest, bagging and decision trees.

3.1 Data

Two sets of examples of human behavior were used for evaluating the performance of the posture classifiers. The first one, containing 135 sequences of behavior of three persons, includes examples of standing/walking, lying down, sitting down, and falling. The second set, which contains 775 sequences of behavior of five people, includes the basic behaviors recorded in the first set, examples of several kinds of falls and, based on discussions with physicians, examples of walking and lying of people with different health problems, such as Parkinson's disease, hemiplegia etc.

The recordings of human behavior used in these experiments were made with the use of the Smart infrared motion capture system [9], because the CONFIDENCE hardware is not available yet. In these recordings, the location of twelve tags was measured, one on each shoulder, hip, knee, ankle, elbow and wrist. The location of a virtual tag on the neck was computed as the middle point between the shoulders due to difficulties in attaching a tag there and tracking it during forward falls. The coordinates of the tags were sampled at a frequency of 60 Hz. This data was processed in order to bring it in a form analogous to the data we anticipate to obtain with the CONFIDENCE hardware. Because the CONFIDENCE hardware will use the ultra-wideband technology, the same technology on which the Ubisense system [10] is based, we expect its properties to correspond to the ones of Ubisense. For this reason, we applied two transformations on the Smart data in order to make it analogous to the data of Ubisense. First, the sampling frequency was reduced to 10 Hz. Then, Gaussian noise with standard deviation of 4.36 cm horizontally and 5.44 cm vertically was added to the data. The values of the standard deviation of the noise in the Ubisense system were obtained experimentally.

Table 2: Distribution of classes per phase

Posture	First phase	Second phase
Standing	1544	39070
Sitting	733	5368
Lying	1773	5337
Falling	689	2229
Moving downwards	1696	5044
Moving upwards	0	421
On all fours	0	2183

The experiments presented in this paper use data only about the location of the neck and the ankles. The posture classification rule engine uses this data as input (Section 2) and for the purpose of comparison the machine learning models were induced on the same data. Each instance, representing the user's body configuration at given moment, contains the following attributes:

- Z-coordinate of the neck and ankles
- Absolute neck-ankle distance
- Neck-ankle distance in Z-direction and its projection on the XY-plane
- Absolute velocity and velocity in Z-direction of the neck and the ankles.

The attributes were computed as averages over half a second overlapping sequences of behavior. The distribution of the classes of interest is given in Table 2.

3.2 Results

The performance of the rule engine was compared with the performance of four machine learning techniques – support vector machines, random forest, bagging and decision trees (Table 3). The machine learning techniques were evaluated with the use of Weka [11]. The support vector machines, random forest and bagging classifier were induced with the default Weka setting. The decision tree classifier was induced with the minimal number of instances per leaf set to 2% of the training dataset. The performance of the machine learning techniques was evaluated with 10-fold cross validation on the data from both phases and with three separate training and test set scenarios. In the first and second separate training and test set scenario, the classifier was induced on data from one phase of recordings and tested on the other. In the third scenario, the classifier was trained on recordings of two persons from both phases and tested on recordings of a third person. The performance of the rule engine is presented with its accuracy on the test dataset for each separate training and test set scenario.

Examination of the classification accuracy of the machine learning techniques in the different evaluation scenarios suggests a certain degree of over-fitting. The accuracy of these classifiers is highest when evaluated with 10-fold cross validation. The random selection of training and test dataset in 10-fold cross validation leads data about the behavior of concrete person in a concrete phase to be present in both the training and test dataset. Therefore, over-fitting is most likely to be present in this evaluation scenario. The classification accuracy falls when the classifiers are induced on data about two persons and tested on data of a third person. In this case, the training dataset does not contain data about the behavior of the person on which the model is tested. However, since all persons were instructed to behave in the same way in both phases of recordings and they were able to observe and copy each other, the models induced in this evaluation scenario are likely over-fitted to this particular behavior of the persons. The most significant drop in accuracy happens when the classifiers are induced on one phase of recordings and tested on the other. In this case, the training and test dataset contain different behavior and there are persons for which recordings were only made in the second phase. The fall of classification accuracy in this scenario confirms that the models induced with machine learning get over-fitted to the persons and behavior present in the training dataset.

As seen in Table 3, the classification accuracy of

Table 3: Comparison of the classification accuracy among the rule engine and four machine learning techniques

Evaluation		Support vector machines	Random forest	Bagging	Decision trees	Rule engine
Training dataset	Test dataset					
First phase	Second phase	84.68%	74.01%	74.76%	79.92%	90.91%
Second phase	First phase	77.96%	79.63%	79.25%	70.78%	81.76%
Two persons	Third person	86.89%	85.99%	86.52%	87.24%	89.42%
10-fold cross validation		89.51%	95.85%	94.91%	87.66%	

the decision trees techniques in the different evaluation scenarios is more than 10 percentage points lower than the classification accuracy of the rule engine when the classifier is trained on one phase of recordings and tested on the other. There was no significant difference in accuracy between the decision tree and the rule engine when the trees were trained on the data of two persons and tested on a third. This shows that the incorporation of expert knowledge in the knowledge obtained by inducing decision trees increased the generality of the posture classifier.

The accuracy of the other three machine learning techniques is also smaller when training is done on one phase of recordings and testing on the other. The difference is especially significant for the random forest and bagging classifier, where the drop in accuracy is more than 15% when the classifier is induced on the data from the first phase of recordings and tested on the data from the second phase. The difference in classification accuracy between the machine learning techniques and the rule engine is not significant when the classifiers are trained on data about two persons and tested on a third person. Nevertheless, the higher classification accuracy of the rule engine still suggests that the incorporation of common sense improved the generality of the classifier.

4 Conclusion

This paper presented a rule engine for classifying posture which was developed by combining machine learning and expert knowledge. In particular, the architecture of the rule engine and the procedure for rule extraction was described.

The approach achieved 5 percentage points increase in classification accuracy compared to support vector machines with only three tags on user's body. Compared to the decision trees, the improvement was a very significant 11 percentage points. Combining machine learning technologies with expert knowledge proved to be successful at least in the experiments performed. We expect that the constructed posture classification rule engine will be even more robust than the machine learning models in real-life circumstances.

Having identified the relevant attributes for each of the postures of interest, the next step for development of the posture classification rule engine is to automate the determination of the limits in the conditions of the rules. What is more, since they depend on the body configuration of the user, the automation must encompass adaptation to the user.

The procedure for enriching knowledge extracted from decision trees with expert knowledge is not limited to posture classification. It can be applied to any area in which representative training dataset is not available or is difficult to obtain.

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.

References

[1] CONFIDENCE, http://www.confidence-eu.org/. [2009-07-06]

[2] N. González, I. Vélez, CONFIDENCE: Support To Older People's Independent Living, *Proceedings of Psykologia* (2008).

[3] C. Quemada, I. Val, M. Pietrzyk, T. Von der Grün, N. González, I. Vélez, A Smart Care System for Elderly People to Support Independent Living: CONFIDENCE, Smart *Systems Integration* (2009).

[4] M. Luštrek, B. Kaluža, Fall detection and activity recognition with machine learning, *Informatica 33 2* (2009).

[5] M. Luštrek, M. Gams, I.Vélez, Posture and movement monitoring for ambient assisted living, IST-Africa 2009, 6-8 May 2009, Kampala (Uganda).

[6] M. Luštrek, M. Gams, Posture And Movement Recognition From Locations Of Body Tags, Ami-08 (European Conference on Ambient Intelligence), 19-22 November 2008, Nürberg (Germany).

[7] B. Kaluža, M. Luštrek, Fall Detection and Activity Recognition Methods for the Confidence Project: A Survey, IS 2009, 12-16 October 2009, Ljubljana (Slovenia), *Proceedings of the 11th International Multiconference Information Society-IS 2008*, Volume A.

[8] M. Ožek, M. Gams, Use of Data Mining Techniques for Process Analysis on Small Databases, *Dynamic and Advanced Data Mining for Progressing Technological Development*, IGI Global, USA (accepted).

[9] eMotion, Smart motion capture system, http://www.emotion3d.com/smart/smart.html. [2009-07-06]

[10] Ubisense, http://www.ubisense.net/. [2009-07-06]

[11] I. H. Witten, E. Frank, *Data Mining: Practical machine learning tools and techniques*, 2nd Edition, Morgan Kaufmann, San Francisco, USA, 2005.