Journal of Medical Systems manuscript No.
(will be inserted by the editor)

Activity recognition for diabetic patients using a smartphone

Božidara Cvetković · Vito Janko · Alfonso E. Romero · Özgür Kafalı · Kostas Stathis · Mitja Luštrek

Received: date / Accepted: date

Abstract Diabetes is a disease that has to be managed through appropriate lifestyle. Technology can help with this, particularly when it is designed so that it does not impose an additional burden on the patient. This paper presents an approach that combines machine-learning and symbolic reasoning to recognise high-level lifestyle activities using sensor data obtained primarily from the patient's smartphone. We compare five methods for machine-learning which differ in the amount of manually labelled data by the user, to investigate the trade-off between the labelling effort and recognition accuracy. In an evaluation on real-life data, the highest accuracy of 83.4 % was achieved by the MCAT method, which is capable of gradually adapting to each user.

Keywords Activity recognition · Smartphone · Lifestyle · Diabetes

B. Cvetković, V. Janko, M. Luštrek
Jožef Stefan Institue
Jamova cesta 39, Slovenia
Tel.: +386-1-477-3498
E-mail: {boza.cvetkovic, vito.janko, mitja.lustrek}@ijs.si

B. Cvetković, V. Janko, M. Luštrek
Jožef Stefan International Postgraduate School
Jamova cesta 39, Slovenia

Özgür Kafalı
North Carolina State University
Raleigh, NC 27695-8206, USA
E-mail: rkafali@ncsu.edu

Alfonso E. Romero, Kostas Stathis
Royal Holloway, University of London
Egham, TW20 0EX, UK
E-mail: {alfonso.romero, kostas.stathis}@rhul.ac.uk

Introduction

Currently around 415 million people suffer from diabetes, and by 2040 their number is expected to increase to 642 million [3]. Diabetes is a metabolic disease in which the body's cells do not respond properly to insulin or the insulin production is inadequate, so glucose is not removed from the blood to be used for energy. Since diabetes has no cure, it has to be managed through medication and appropriate lifestyle. Key lifestyle activities for diabetic patients are eating and exercise – eating causes glucose in the blood to rise while exercise speeds up its absorption.

Technology for activity monitoring can help the patients better manage their lifestyle, and provide their physicians an insight into the patients' life. In the COMMODITY12 project [2], we developed a personal health system for diabetes intended both for the patients and their physicians. In order not to burden the patients with unnecessary devices, it primarily relies on a smartphone. The phone can be augmented with an ECG monitor, which was introduced for the management of cardiovascular co-morbidities [13] and diagnosis [15], but is also used for activity recognition.

Activity recognition is a common task in telemonitoring, because physical activities characterise the patients' lifestyle and provide the context for more direct health-related observations. For example, in diabetes blood glucose needs to be interpreted in the context of the two most important activities which influence the blood glucose level: eating and exercise. The recognition of other activities gives an insight into the patients' lifestyle which makes it possible to give more appropriate lifestyle advice and monitor the progress of the disease and its complications.

Activity recognition is usually done by utilizing environmental or wearable sensors or even both [17]. This paper focuses on wearable sensors. Most approaches for activity recognition with smartphone sensors and other wearable devices tackle low-level activities such as Walking, Sitting and Lying. In contrast, this paper describes the recognition of high-level activities such as Work, Eating, and Home activities.

Low-level activities are commonly recognized with machine learning from a sliding window passed over a stream of acceleration data [4][5]. Dernbach et al. [8] used the low-level approach for high-level activities and reached the accuracy of barely 50 %. Lee & Cho [18] applied hierarchical hidden Markov models to accelerometer data to first determine low-level activities, and from those high-level activities (Shopping, Taking_bus, Walking). They reached the precision of around 80 %, but their set of activities was very limited and the users carried the phone in their hand.

Using other smartphone sensors, such as the GPS, improves activity recognition. This way, Lin [19] classified Work, Sleep, Leisure, Visit, Driving and Other with conditional random fields. He achieved the accuracy of 86 %. Wang et al. [21] additionally used the microphone and light sensor. They determined the users' state (Working, Home_talking, Place_speech etc.) with a rule-based system. They achieved the accuracy of around 90 %, but the users' home and office wi-fi names were known to the system in advance, and some states were defined expressly in terms of ambient sounds.

Helal et al. [11] developed a platform specifically for monitoring diabetic patients. They used hidden Markov models to recognize activities such as Washing_hands, Cooking and Eating_oatmeal. They relied on ambient sensors, restricting their platform to the patients' homes. Amft et al. [1] detected eating from video, and recognised different types of eating by analysing body movement while eating. However, this was done with fairly intrusive inertial, EMG and other sensors.

Unlike in most related work, the system described in this paper attempts to recognize all the users' activities, including ambiguous ones (e.g., cycling can be exercise, transport or a part of shopping), and was tested in a completely natural setting. The core methods are based on machine learning. We investigated various options differing in the number of models used, and the amount of data labelled by the person using the system. The recognized activities are refined by rules encoded in Event Calculus [16] and optimized for speed and real-time purposes.

Activity recognition approach

Our activity recognition approach is composed of three main steps: feature extraction, core activity recognition, and refinement of the recognized activities. Feature extraction transforms raw sensor data into features that can be used in the next two steps. The core activity recognition mainly uses machine-learning to recognize high-level lifestyle activities. Finally, we refine the recognized activities with symbolic reasoning.

Features for activity recognition

The features are extracted from smartphone sensors and optionally from an accelerometer-equipped chest-worn ECG monitor. They are computed over one-minute windows and belong to five groups: sound, location, acceleration, heart-rate and respiration-rate.

Sound features are extracted from the ambient sound recorded with the smartphone's microphone using the jAudio library [12]. We record 100 ms of sound out of each second in a minute (to preserve the users' privacy) and further split it to 20 ms sub-windows. The computed features are the average spectral-centroid, zero-crossing, mel-frequency-cepstral-coefficient, method-of-moments values and linear-predictive-coding for each one-minute window [6].

Location features are extracted from two sources: (i) the smartphone's GPS receiver, from which we take the geographical coordinates, and from (ii) the wi-fi module, from which we take information about the visible access points and signal strength. The features from both sources are clustered into 20 clusters each with the hierarchical clustering as implemented in the Weka machine-learning suite [10]. The clusters are afterwards assigned one of three semantic locations – residence, work and elsewhere – with the following algorithm:

– The importance of each cluster is computed as the fraction of time the user spends in the cluster in a day (heatmap in Fig. 1).
– Clusters with the highest free-day importance, that also have workday importance above 0, belong to the location residence (blue semantic location in Fig. 1).
– Clusters with highest workday importance above 0, which are never visited on free days, and aggregate into 50% of the time, belong to the location work.
– All other clusters belong to the location elsewhere.

Each wi-fi cluster is assigned the correct semantic location with the accuracy of 84 % and each GPS cluster with 81 %. In Fig. 2 we present a GPS plot (latitude,

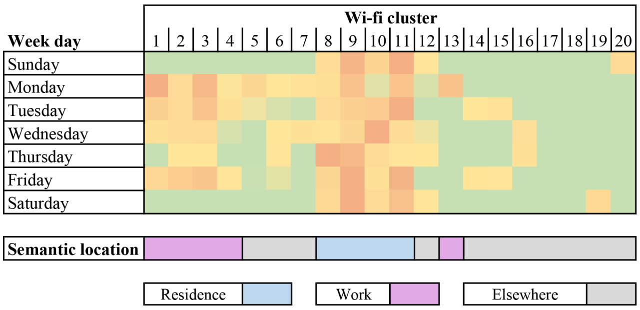

Fig. 1: Heatmap for wi-fi clusters and the transformation into semantic locations for an example user.

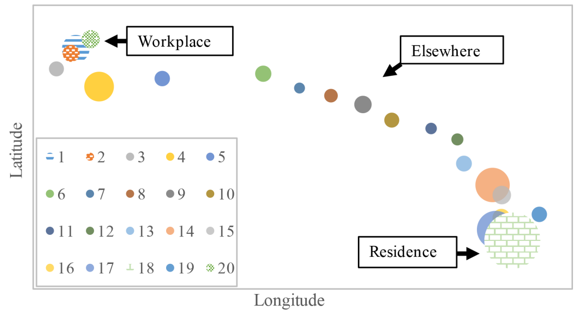

Fig. 2: All 20 GPS clusters and their corresponding semantic locations for an example user.

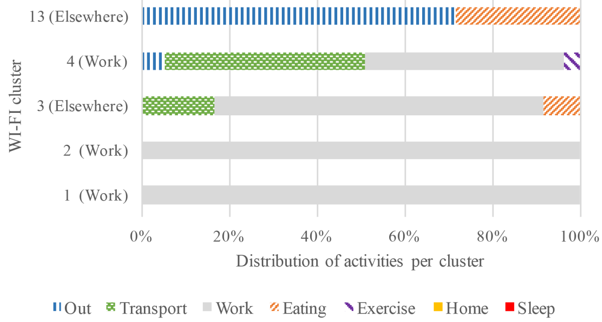

Fig. 3: Wi-fi clusters composing GPS clusters 1, 2 and 20 (GPS work semantic location). Each wi-fi cluster is labelled with the recognised wi-fi semantic location and presented with the distribution of activities.

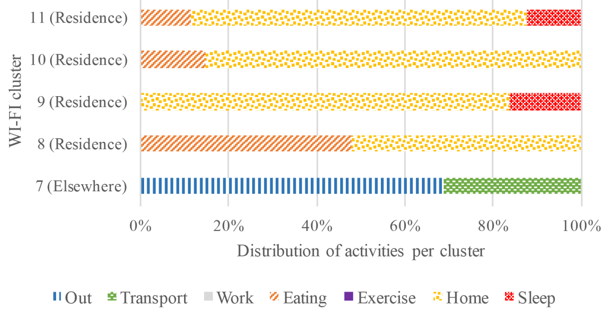

Fig. 4: Wi-fi clusters composing the GPS cluster 18 (GPS residence semantic location). Each wi-fi cluster is labelled with the recognised wi-fi semantic location and presented with the distribution of activities.

longitude) with GPS clusters (the size of the circle represents the size of the cluster) and their corresponding semantic locations. Since GPS can only be used outdoors due to its poor signal indoors, the wi-fi signal clustering is needed for more fine-grained location indoors. Fig. 3 shows workplace clusters from Fig. 2 (GPS clusters 1, 2 and 20) and Fig. 4 the residence cluster from Fig. 2 (GPS cluster 18) as a composition of wi-fi clusters, assigned semantic location and the distribution of activities per cluster. We can observe that wi-fi clusters enable detection of more fine-grained semantic locations (in brackets after each cluster number).

GPS work location (GPS cluster 1, 2 and 20) is composed of five wi-fi clusters (Fig. 3): wi-fi-clusters 1 and 2 are clean work clusters, wi-fi cluster 2 is a mixed cluster corresponding to a cafeteria close to the office, the wi-fi cluster 4 is probably assigned to the parking lot close to the office (Exercise activity represents the running to the car), and wi-fi cluster 13 represents the restaurant near the workplace.

GPS residence location (GPS cluster 18) is composed of five wi-fi clusters (Fig. 4): wi-fi cluster 7 represents the outdoors of the residence, while the remaining four may represent different parts of the house. Wi-fi cluster 8 and 10 probably represent the dining room, wi-fi cluster 9 the bedroom and wi-fi cluster 11 the living room.

Additional location features are the wi-fi availability determined by the strength of the wi-fi signal, GPS availability determined by the visibility of satellites, velocity, and location category according to the Foursquare location API [9].

Acceleration features are extracted from the smartphone's and/or ECG monitor's accelerometer. They are the user's most common low-level activity, the percentage of each low-level activity within the one-minute window, the user's average estimated expended energy and the percentage of the estimated low, medium and vigorous intensity of the activity. They are computed with our recent method [5] that can seamlessly use the phone, ECG monitor or both.

Heart-rate features and *respiration-rate features* are extracted from the ECG monitor if present. The features are the (i) minimum, (ii) maximum and (iii) average heart-rate or respiration-rate within each one-minute window.

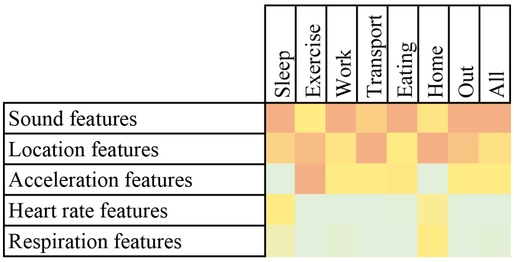

Fig. 5: Correlation heatmap between groups of features and activities.

The last three categories of features use the ECG monitor. The quality of the acceleration features is not much degraded without the ECG monitor [5], and according to the heatmap presented in Fig. 5, the heart-rate and respiration features are not very important, so the ECG monitor is not essential.

Core methods for activity recognition

Since each person performs the daily activities differently, our activity-recognition system can benefit from labelled activities of each user. However, since labelling represents a burden for the user, our goal was to investigate the trade-off between the amount of labelled data and the recognition accuracy. We developed five methods, some of which use a single model trained with machine learning (Fig. 6) and some multiple models (Fig. 7):

– *Person-dependent method (PD, Fig. 6a)* is a single-model method, where the model (PD model) for each user is trained on that user's data only (PD dataset). This requires the user to label all the activities to be recognised during the training phase – one week in our experiments. This can be fairly burdensome, and the quality of the recognition strongly depends on the user's conscientiousness.
– *Person-independent method (PI, Fig. 6b)* is a single-model method, where the model (PI model) for each user is trained on data of people other than the user (PI dataset). This approach does not require any labelling from the user.
– *PI method with person-specific data (PIA, Fig. 6c)* is a single-model method, where the PI dataset is augmented with user-specific data of Exercise and Eating (these two activities are the most relevant for diabetic patients, but any activities could be included in principle). This approach requires the user to label Eating and Exercise during the training phase, after which an augmented model is trained (PIA model).

Table 1: Comparison of the amount of manually labelled data by the user in the core methods for activity recognition.

	Labelled data by the user	
Method	**Activities**	**Duration**
PD	All	1 week
PI	None	None
PIA	Eating, Exercise	1 week
PIAH	Eating, Exercise	1 week
MCAT	Eating, Exercise, Home	First occurrence

– *PI and person-specific model combined with heuristics method (PIAH, Fig. 7a)* is a two-model method. The PI model is trained on the PI dataset as in the PI method. The user-specific model is trained on data of Exercise and Eating labelled by the user, requiring the same effort as the PIA method. The outputs of the two models are merged with a heuristic method into the final activity. The reader is referred to [6] for more details.
– *Multi-Classifier Adaptive Training (MCAT, Fig. 7b)* is a semi-supervised learning method that can adapt to a user while the system is in use. It is composed of two domain-dependant models and two domain-independent meta-models. The domain-dependant models are the PI model and a small user-specific model. The user-specific model is trained on only the first occurrences of Eating, Exercise and Home activity, requiring minimal labelling effort from the user (again, any activities could be used in principle). The outputs of the two models are merged with a meta-model that selects which of them to trust. The MCAT algorithm also utilises a meta-model that decides whether each instance should be included in the PI dataset or not. The PI model is then periodically retrained, becoming ever more personalised. While the MCAT method does not need much user-labelled data, it does need a fairly large amount of unlabelled data to gradually adapt to the user. The reader is referred to our earlier work on a different problem [7] for more details.

All five methods are summarised in Table 1 according to the labelling effort needed from the user.

Refinement of recognised activities

Activity recognition using a mobile phone is error prone because mobile sensors are noisy. Sensor noise fed in classification algorithms propagate errors in predictions and as a result an activity being recognised can be incorrect. On top of this, classification algorithms are

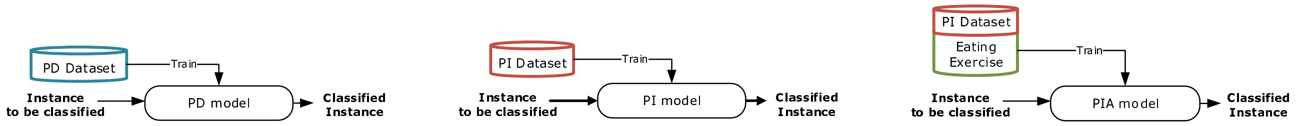

(a) Person dependent method (b) Person independent method (c) PI with person-specific data included method

Fig. 6: Single-model methods.

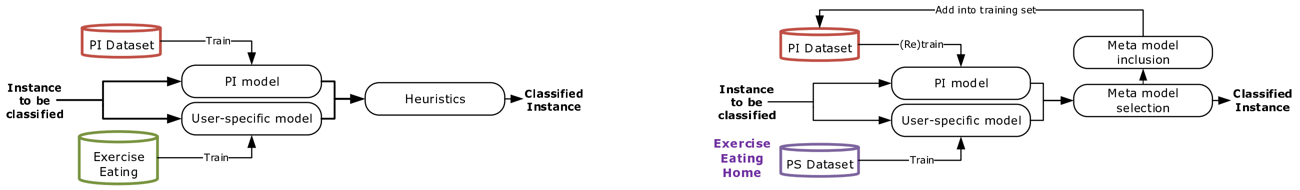

(a) PI and person-specific model combined with heuristics method (b) Multi-Classifier Adaptive Training method

Fig. 7: Multi-model methods.

oblivious to common sense knowledge, thus it is possible for example to confuse that a person is eating while the person is actually exercising, even if it is commonly known that eating and exercising to be incompatible.

To refine confused predictions we study how symbolic rules can be used with classification algorithms. The idea is to refine the classification models and to act as interfaces with techniques used to monitor a diabetic patient and reason about his condition as in the agent-oriented system discussed in [13]. More specifically, Kafalı et al [14] show how activity recognition from sensors can be reasoned upon with symbolic rules to arrive at more complex conclusions, for example, that while a person was walking home from work has fallen because he had a hypoglycemic attack. This framework is based on a specific representation of activities where events (or observations) that happen at different times start, suspend, resume and achieve activities according to a user's goals specified at the mobile phone's settings.

We combine the Kafalı et al framework together with predictions provided by a classifier to eliminate misclassified activities using symbolic rules of the form: *Conclusion* ← *Conditions* where the *Conclusion* is a predicate that holds to be true if the set of predicates in the *Conditions* hold to be true. A predicate name must be a constant, while each argument within a predicate can either be a constant or a variable. Constants are represented by identifiers starting with a lower-case letter and variables by identifiers starting with an uppercase letter. A constant stands for a specific entity, and different constants stand for different entities. A variable can stand for any entity, and different variables can stand for the same entity. A single underscore '_' denotes an anonymous variable and means 'any term'.

Removal of orphan activities (R1). If the classifier predicts an activity A at a stage identified by time T+1, an activity B at the previous stage identified by time T, and the activity A at a previous stage identified by T-1, then the activity B is replaced with activity A at the stage identified by T. Such a generic rule eliminates predictions such as a person being at home, then being at work, then being back at home in three consecutive and very close to each other times, and replaces them with the person being at home throughout these times.

```
possible_at(activity(A, P)=active, T) ←
  predicted(A, _, T+1),
  predicted(B, _, T),
  A ≠ B,
  predicted(A, _, T-1).
```

As an example of how this rule can be applied, we show the following set of predicted activities taking from the dataset: (12:01:48, Eating), (12:03:00, Transport), (12:04:12, Transport). In that case, the second prediction will be changed into Eating, as it seems unlikely that there is one minute of Transport.

Learning person-dependent activity intervals (R2) captures whether an activity is valid at a given time. For instance, we know that an average person usually eats at least three times a day. If we know that a person eats at specific times, we validate the predicted eating activity. Moreover, we learn such valid activity intervals for each activity using the training datasets available for each individual person.

```
possible_at(activity(A, P)=active, T) ←
  member(A, [sleep, work, eating]),
  valid_interval(P, A, Interval),
  member(Ts-Te, Interval),
  Te ≥ T, T ≥ Ts.

valid_interval(p1, sleep, [30-730]).
valid_interval(p1, work,
      [915-1200,1300-1730]).
valid_interval(p1, eating,
      [845-900,1200-1230,1800-1830]).
```

For a certain person of the dataset, for example, the maximal sleep interval across all the training data is [00:16:12, 10:12:12]. A prediction of Sleep found at 20:40:48 will be changed into the second more likely prediction (Home, in this case).

We apply the rules in the following order: R1 then R2. R1 applies to all activities, and R2 applies to Sleep, Work and Eating.

Experiments

The activity-recognition approach was evaluated on the recordings of nine volunteers who wore the smartphone and the ECG monitor for two weeks each. All the methods used the first week of recordings for training and the second week for testing. Person-independent methods (PI, PIA, PIAH and MCAT) were evaluated with the leave-one-person-out approach, meaning that models were trained on the data of eight people (first week) and tested on the ninth (second week), repeated once for each person. The first week of each test person was used to determine the semantic locations (residence, work and elsewhere), and to train user-specific classifiers where applicable (see Table 1).

The models in the core activity-recognition are ensembles of base models trained with SVM, J48, Random Forest, JRip, AdaBoost and Bagging algorithms with default parameters as implemented in the Weka machine-learning suite [10]. The exception are the user-specific and two meta-classifiers in the MCAT method, the former of which is trained with the Random Forest algorithm and the latter two with the J48 algorithm.

Dataset

The nine healthy volunteers (eight male, one female, aged from 24 to 36) were asked to carry the smartphone with the data-collection application as much as possible for two weeks, in any pocket they wanted (or in a bag), and to wear the ECG monitor each day until the battery ran out. They were asked to lead their life as usual. While some of them had fairly regular daily routines, others had unusual eating patterns, were staying at a different place during the week and weekend, took trips etc., so the resulting dataset is quite challenging for activity recognition. On average, we collected 7.5 hours of recordings per day with the ECG monitor and 11 hours with the phone.

The smartphone application was collecting the sensor data and was used for labelling the activities: Home-chores, Home-leisure, Food preparation, Eating, Exercise, Work, Out-errands, Out-leisure, Sleep and Transport. We later merged Home-chores, Home-leisure and Food preparation into Home, and Out-Errands and Out-leisure into Out, since these activities proved impossible to distinguish.

Results

Table 2 shows the results of all the evaluated methods in terms of micro-averaged classification accuracy and f-score, both before and after the refinement with symbolic reasoning. In Table 3 we present the f-scores per activity after the refinement. The default model, which serves as a baseline for comparison, always outputs the most common class, which is the work activity.

We first evaluated the two simplest methods: PD, which requires the user to label all the activities during the training week, and PI, which requires no user labelling. The advantage of the PD method is that it can adapt to the user, but it has less training data available. The advantage of the PI method is more training data (from eight people vs. one), but it is not adapted to the user. The results of both methods proved similar, so apparently their advantages and disadvantages balanced out. The PI method was more successful at recognising the Transport and Out activities, which reflects their generality, while the PD method was more successful at other activities, particularly Eating and Exercise, which suggests they are more user-specific.

We continued with the evaluation of methods combining user-independent and user-specific data and models. The PIA method, which simply adds user-specific Exercise and Eating data to its model, significantly improved the recognition of Exercise, but had less success with Eating and was the least accurate overall. The PIAH method, which heuristically merges a person-independent and person-specific model, proved almost identical to the PI method overall. It did yield a higher accuracy for four people (81 % – 85 %), but a lower one for the rest (50 % – 76 %). This indicates that some recorded people had substantially different training and test weeks.

At the end we evaluated the MCAT method [7], for which the user was required to label a single Eating, Ex-

Table 2: Average classification accuracy (%) and micro-averaged f-score for the evaluated core methods and after correction with symbolic reasoning.

	Core methods		Refining method	
Methods	Accuracy	F-score	Accuracy	F-score
Default	35.9	0.35	35.9	0.35
PD	72.9	0.73	74.3	0.73
PI	71.3	0.71	73.7	0.72
PIA	66.4	0.66	69.9	0.67
PIAH	72.0	0.72	73.7	0.72
MCAT	81.8	0.82	83.4	0.82

Table 3: Average f-score for the activities after machine-learning + symbolic approach approach.

	Methods				
Activity	PD	PI	PIA	PIAH	MCAT
Sleep	0.80	0.72	0.71	0.72	0.80
Exercise	0.62	0.40	0.58	0.40	0.91
Work	0.88	0.85	0.81	0.85	0.90
Transport	0.58	0.69	0.71	0.69	0.78
Eating	0.34	0.22	0.24	0.22	0.68
Home	0.81	0.75	0.70	0.75	0.83
Out	0.35	0.51	0.51	0.51	0.69

ercise and Home activity. The MCAT method gradually adapts the PI model to a current user by retraining it on automatically labelled and selected data. It was retrained three times in our experiments. The adaptation improved the recognition of all the activities, not only those for which it had user-specific data. The confusion matrix of the MCAT method with the symbolic refinement is presented in Table 4. Most of the activities were miss-classified as the Home because they occurred at the residence location (the location being the most important feature for activity recognition). For example, most of the recorded people exercised at their residence, and all of them ate and slept there. The miss-classification of Transport occurred during transitions from/to Home, Out and Work. Eating was confused not only with Home, but also Work, Out and even Transport. This happened because people were eating in the same locations and postures as doing home activities (e.g., eating at the kitchen table or on the sofa where they also read and watch TV), working (e.g., eating at their office desk), eating while shopping or commuting.

Conclusion

In this paper we presented an approach that combines machine learning and symbolic reasoning to recognise high-level lifestyle activities of diabetic patients using sensor data obtained primarily from the patients' smartphones. Our motivation was that diabetes has to be managed through appropriate lifestyle – particularly eating and exercise, which directly affect blood glucose. Lifestyle is important for other diseases as well, so our work has applications beyond diabetes. The main challenge was that patients' lifestyles and consequently ways they perform activities are very diverse, so activity recognition should be adapted to each individual. But since labelling represents a burden for the user, we focused on investigating the trade-offs between the amount of labelled data and the recognition accuracy.

Our first results showed that the two extreme cases – relying only on user-labelled data and completely eschewing user-labelled data – are comparable. This suggested that the strengths of both methods should be combined. We did this in increasingly sophisticated fashion: by simply merging general and user-specific training data (PIA method), by heuristically merging the outputs of general and user-specific classifiers (PIAH method), and by utilising semi-supervised learning approach (MCAT method). The last method required the least labelled user-specific data, but nevertheless produced the best results. This is because it could take advantage of unlabelled data, which is easily obtained in large quantities, since it requires no effort from the user other then carrying a smartphone with an activity-monitoring application. Since not all knowledge about activities is captured by phone sensors, the machine-learning methods benefited from symbolic refinement based on human understanding of the problem, albeit the benefit was not very large.

We believe that the accuracy of the MCAT method is sufficient to be of real benefit for diabetic patients. The overall accuracy of 83.4 % and f-score of 0.82 are comparable to related work, even to examples that tackle easier problems. The recognition of Exercise with the f-score of 0.91 is even better. Only the recognition of Eating with the f-score of 0.68 is somewhat lacking. Since unobtrusive monitoring with smartphone sensors cannot possibly recognise the quantity of food eaten, it can only prompt the user to input the quantity or act (e.g., inject insulin) according to his/her own judgement anyway. In this capacity, even a somewhat lower accuracy is useful. The detection of Sleep can be used to recognise sleep problems, which are common in diabetes. The overall activity recognition can be used to recognise diabetes fatigue and depression by monitoring deviations from the daily routine and the degree of activity. The recognition of Home, Work and Transport activities enables quality advice about exercise (e.g., remind the users of their daily exercise during Home activity rather than work or suggest commuting on foot or by bicycle).

Table 4: Confusion matrix of the MCAT method refined with the symbolic reasoning. The brackets contain the contribution of the refining method (green is increase in accuracy, pink is decrease and no colour is no change).

	PREDICTED						
TRUE	Sleep	Exercise	Work	Transport	Eating	Home	Out
Sleep	661 (-2)	0 (0)	0 (0)	0 (0)	0 (0)	292 (2)	0 (0)
Exercise	0 (0)	859 (11)	3 (-1)	5 (-3)	8 (-5)	53 (-1)	2 (-1)
Work	31 (0)	18 (0)	11322 (68)	173 (-4)	110 (-6)	922(-59)	49 (1)
Transport	9 (-2)	13 (-1)	244 (6)	2927 (11)	18 (-3)	235 (-2)	142 (-9)
Eating	1 (0)	1 (0)	50 (-3)	10 (-2)	1329 (15)	497 (-11)	200 (1)
Home	164 (-67)	42 (2)	495 (-11)	71 (0)	138 (-18)	9377 (115)	60 (-21)
Out	0 (0)	20 (-2)	390 (-1)	757 (-5)	197 (-12)	822 (-10)	2884 (30)

In the future, we plan to use the activity recognition methods described in this paper for just such lifestyle advice, as well as advanced diseases monitoring. We will also improve the recognition of Eating by replacing the ECG monitor with a wrist-worn device. Such devices are also preferable because they are more comfortable and can be worn at all times. Another area for future work is sound processing, since the sound is a general feature with the potential to improve the recognition of eating as well as other activities.

Acknowledgements This paper significantly extends the work of Luštrek et al. [20], which was presented in the PHSCD 2015 Workshop. The work was partially supported by the EU FP7 project COMMODITY_{12} (www.commodity12.eu).

References

1. Amft, O., Ambient, on-body, and implantable monitoring technologies to assess dietary behavior. Handbook of Behavior, Food and Nutrition. Springer, 2011.
2. Commodity12, http://www.commodity12.eu/. Accessed 15 February 2016.
3. Diabetes Atlas, http://www.diabetesatlas.org/. Accessed 15 February 2016.
4. Cvetković, B., Milić, R. and Luštrek, M., Estimating Energy Expenditure with Multiple Models using Different Wearable Sensors. IEEE Journal of Biomedical and Health Informatics, accepted for publication in July 2016.
5. Cvetković, B., Janko, V., and Luštrek, M., Demo abstract: Activity recognition and human energy expenditure estimation with a smartphone. In PerCom 2015, 23–27 March 2015, St. Louis, USA, 2015.
6. Cvetković, B., Mirchevska, V., Janko, V., and Luštrek, M., Recognition of high-level activities with a smartphone. In Adjunct Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2015 ACM International Symposium on Wearable Computers (UbiComp/ISWC'15 Adjunct), 1453–1461, 2015.
7. Cvetković, B., Kaluža, B., Gams, M., and Luštrek, M., Adapting activity recognition to a person with Multi-Classifier Adaptive Training. Journal of Ambient Intelligence and Smart Environments 7(2), 171–185, 2015.
8. Dernbach, S., Das, B., Krishnan, N. C., Thomas, B. L., and Cook, D. J., Simple and complex activity recognition through smart phones. In 8th International Conference on Intelligent Environments, pp. 214–221, 2012.
9. Foursquare API. https://developer.foursquare.com/. Accessed 15 February 2016.
10. Hall, M., Frank, E., Holmes, G., Pfahringer, B., Reutemann, P., and Witten, I. H., The WEKA data mining software: An update. SIGKDD Explorations 11(1), 10–18, 2009. doi:http://dx.doi.org/10.1145/1656274.1656278
11. Helal, A., Cook, D. J., and Schmalz, M., Smart home-based health platform for behavioral monitoring and alteration of diabetes patients. Journal of Diabetes Science and Technology, 3(1), 141–148, 2009.
12. jAudio library. http://jaudio.sourceforge.net/. Accessed 15 February 2016.
13. Kafalı, Ö., Bromuri, S., Sindlar, M., Van der Weide, T., Pelaez, E.A., Schaechtle, U., Alves, B., Zufferey, D., Rodriguez-Villegas, E., Schumacher, M.I., and Stathis, K., COMMODITY_{12}: A smart e-health environment for diabetes management. Journal of Ambient Intelligence and Smart Environments. 5(5) 479–502, 2013. doi:10.3233AIS-130220.
14. Kafalı, Ö., Alfonso, E.R., and Stathis, K., Activity Recognition for an Agent-oriented Personal Health System. Principles and Practice of Multi-Agent Systems – 17th International Conference, PRIMA 2014, Springer, 2014.
15. Kafalı, Ö., Schaechtle, U., and Stathis, K., Hydra: A hybrid diagnosis and monitoring architecture for diabetes. 16th IEEE International Conference on e-Health Networking, Applications and Services, Healthcom 2014, pp. 531–536, Natal-RN, Brazil, October 15–18, 2014.
16. Kowalski, R., and Sergot, M., A logic-based calculus of events. New Generation Computing, 4(1), 67–95, 1986.
17. Lara, O.D., and Labrador, M. A.:A Survey on Human Activity Recognition using Wearable Sensors. IEEE Communications Surveys & Tutorials, 15(3), 1192–1209, 2013.
18. Lee, Z.S., and Cho, S.B., Activity recognition using hierarchical Hidden Markov Models on a smartphone with 3D accelerometer. Hybrid Artificial Intelligent Systems, Springer, Berlin Heidelberg, pp. 460–467, 2011.
19. Lin, L, Location-Based Activity Recognition. Ph.D. Dissertation. University of Washington, 2006.
20. Luštrek, M., Cvetković, B., Mirchevska, V., Kafalı, Ö., Romero, A., Stathis, K. Recognising lifestyle activities of diabetic patients with a smartphone. Pervasive Health 2015: Workshop on Personal Health Systems for Chronic Diseases (PHSCD 2015).
21. Wang, Y., Lin, J., Annavaram, M., Jacobson, Q.A., Hong, J., Krishnamachari, B., Sadeh, N., A framework of energy efficient mobile sensing for automatic user state recognition. In Proceedings of the 7th International Conference on Mobile Systems, Applications, and Services, ACM, New York, USA, pp. 179–192, 2009.