Monitoring Patients' Lifestyle with a Smartphone and Other Devices Placed Freely on the Body

Mitja Luštrek, Božidara Cvetković, and Vito Janko

Jožef Stefan Institute, Department of Intelligent Systems, Ljubljana, Slovenia
{mitja.lustrek,boza.cvetkovic,vito.janko}@ijs.si

Abstract. Monitoring patients' lifestyle can result in an improved treatment, but it is often not critical enough to warrant dedicated sensors. However, many consumer devices, such as smartphones, contain inertial sensors, which can be used for such monitoring. We propose an approach to activity recognition and human energy-expenditure estimation for diabetes patients that uses a phone and an accelerometer-equipped heart-rate monitor. The approach detects which of the two devices is carried or worn, the orientation of the phone and its location on the body, and adapts the monitoring accordingly. By using this approach, the accuracy of the activity recognition was increased by up to 20 percentage points compared to disregarding the orientation and location of the phone, while the error of the energy-expenditure estimation was decreased.

Keywords: Activity recognition · energy expenditure estimation · smartphone · heart-rate monitor · accelerometer · location · orientation · diabetes

1 Introduction

Monitoring patients' lifestyle can often allow physicians to better advise on disease management and to verify the compliance with the prescribed treatment. Our work is motivated by diabetes treatment, where the information about the patients' activities and physical exertion is quite useful, but not critical enough to warrant dedicated sensors. Fortunately the number of consumer devices that contain inertial sensors is increasing, providing opportunity for lifestyle monitoring without dedicated sensors. The prime example of such a device is the smartphone, while others include heart-rate monitors and activity-monitoring wristbands. However, monitoring activities with such devices is challenging because the users can freely choose which devices to carry at any given time, and they can be carried in various orientations and locations.

Activity recognition is probably the most mature area of lifestyle monitoring, and while it is often done with smartphones, the issue of varying orientation and location is rarely addressed. Thiemjarus [1] normalized the orientation of the phone, increasing the accuracy by 16 percentage points. Martín et al. [2] detected the location of the phone and used different activity-recognition models for different locations, increasing the accuracy by 4–10 percentage points. A few researchers also measured physical

adfa, p. 1, 2011.

exertion or human energy expenditure with smartphones [3], but without addressing the varying orientation and location.

In this paper we present an activity-recognition and energy-expenditure-estimation approach using a smartphone and an accelerometer-equipped heart-rate monitor. It automatically adapts to whichever of the two devices is on the user's body, and both normalizes the orientation of the phone and detects its location.

2 The Proposed Activity-Monitoring Approach

The proposed approach can use any number of sensing devices carried or worn by the user, but in our experiments we used two: a smartphone, which could be placed in any orientation in the breast or trousers pocket or in a bag, and an accelerometer-equipped heart-rate monitor worn on the chest. The activity-monitoring pipeline is shown in Fig. 1. First, we detect whether each of the two devices is present on the user's body. Second, if the phone is present, we detect walking in an orientation- and location-independent fashion. When a 10-second period of walking is detected, we use it to normalize the orientation and detect the location of the phone. When the phone is not carried, the information on the orientation and location is reset until the next period of walking. Third, we invoke the activity-recognition model appropriate for the current device configuration, which uses features based on the acceleration data from one or both devices. And fourth, we invoke the appropriate energy-estimation model, which uses the recognized activity, acceleration features and the heart rate if available.

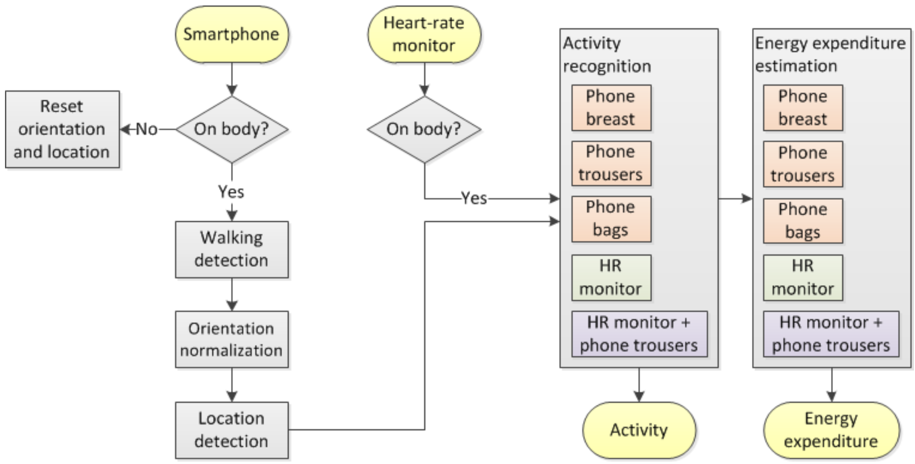

Fig. 1. The activity-monitoring pipeline.

The smartphone is considered **present** on the user's body unless the screen is on, indicating that the phone is in use, there is an ongoing call, or the phone is completely still. The heart-rate monitor can self-report whether it is being worn.

The **normalization of the orientation** is based on the assumption that the average acceleration during walking corresponds to the Earth's gravity. We compute a quaternion matrix representing the rotation between the average orientation of the smartphone during 10 seconds of walking, and the preferred orientation corresponding to the phone's longest side being perfectly aligned with the gravity [4]. This matrix is then applied to acceleration data to normalize it in real time.

The **walking detection**, **location detection** and **activity recognition** are done in essentially the same way, using machine learning. First, the stream of acceleration data is segmented into 2-second windows. Then, a number of features are computed for each window, forming a feature vector. The feature vectors computed from training data are fed into a machine-learning algorithm, which outputs a classification model for one of the three tasks. When new data are obtained, the same features are computed and fed into the model, which outputs the activity or location. The Random Forest algorithm as implemented in the Weka suite [5] was used to train all the classification models, since it outperformed other tested algorithms. The **energy-expenditure estimation** is done similarly as the previous three tasks, except that the window length is 10 seconds and regression models are used instead of classification. The Support Vector Regression algorithm from Weka was used to train the models.

3 Experimental Evaluation

The proposed activity-monitoring approach was evaluated on recordings of 10 volunteers wearing the heart-rate monitor and simultaneously carrying smartphones in all the locations in various orientations. They performed a scenario consisting of everyday activities and exercise, interspersed with periods of walking to normalize the orientation of the phone and detect its location. The activities to be recognized were walking, running, cycling, lying, sitting and standing (the last two were merged when only a device on the chest was available, and the last three when only the phone in the bag was available). The leave-one-person-out cross-validation was used.

The walking detection achieved the classification accuracy of 91.1% with the smartphone alone and 92.2% with both devices. The location detection achieved the classification accuracy of 90.3% with the smartphone alone (the heart-rate monitor was not used). The activity recognition and energy expenditure estimation yielded the results in Table 1. They show that using two devices is worthwhile only when one of them is the smartphone in the trousers pocket. The redundant two-device configurations are marked with a star and not included in the averages. The last column shows

Table 1. The accuracy of the activity recognition [%], and the mean absolute error of the energy expenditure estimation [MET, 1 MET is the energy expended at rest]

		Heart-rate monitor	**Smartphone**				**Comp.**
			Breast	**Trousers**	**Bags**	**Average**	
Activity recognition	One device	81.7	87.9	78.9	79.5	82.1	83.7
	Two devices		83.8*	92.7	60.9*	92.7	92.8
Energy expenditure est.	One device	0.74	0.87	0.92	1.15	0.92	1.00
	Two devices		0.76*	0.72	0.87*	0.72	

comparison approaches: an activity-recognition system using fixed dedicated sensors, trained and tested on a similar dataset as ours, and a state-of-the-art consumer device for energy-expenditure estimation (www.bodymedia.com). Our approach performed similarly to the first despite not using fixed sensors, and better than the second.

To evaluate the gains from the orientation normalization and location detection, we tested the activity monitoring with each of these features turned on or off. Table 2 shows the gains in activity recognition are substantial, while the gains in energy-expenditure estimation are modest and only present if both features are turned on.

Table 2. The accuracy [%] of the activity recognition and the mean absolute error of the energy expenditure estimation [MET] with/without orientation normalization and location detection

	Activity recognition	**Energy expenditure est.**
Without orientation, without location	63.4	0.98
With orientation, without location	68.7	1.17
Without orientation, with location	74.7	1.04
With orientation, with location	83.6	0.92

4 Conclusion

We presented an approach to activity-monitoring of diabetes patients with a smartphone and an accelerometer-equipped heart-rate monitor that can detect whether and how the devices are worn or carried, and adapt the monitoring accordingly. This increased the accuracy of the activity recognition by 20 percentage points and modestly lowered the error of the energy-expenditure estimation. We plan to upgrade the activity recognition with higher-level activities (e.g., work, leisure, eating). Afterwards, the approach will be tested on diabetes patients, whose data is currently being collected. We also plan to apply it to other domains, both in medicine and sports.

References

1. Thiemjarus, S.: A Device-Orientation Independent Method for Activity Recognition. In: International Conference on Body Sensor Networks, pp. 19–23. IEEE, New York (2010)
2. Martín, H., Bernardos, A.M., Iglesias, J., Casar, J.R.: Activity Logging Using Lightweight Classification Techniques in Mobile Devices. Pers. Ubiquit. Comput. 17, 675–695 (2013)
3. Cvetković, B., Kaluža, B., Milić, R., Luštrek, M.: Towards Human Energy Expenditure Estimation Using Smart Phone Inertial Sensors. In: Augusto, J.C., Wichert, R., Collier, R., Keyson, D., Salah, A.A., Tan, A.-H. (eds.) Ambient Intelligence. LNCS, vol. 8309, pp. 94–108. Springer, Heidelberg (2006)
4. Tundo, M.D., Lemaire, E., Baddour, N.: Correcting Smartphone Orientation for Accelerometer-based Analysis. In: MeMeA, pp. 58–62. IEEE, New York, 2013
5. Hall, M., Eibe, F., Holmes, G., Pfahringer, B., Reutemann, P., Witten, I.H.: The WEKA Data Mining Software: An Update. ACM SIGKDD Explorations 11(1), 10–18, 2009
6. Kozina, S., Gjoreski, H., Gams, M., Luštrek, M.: Three-layer Activity Recognition Combining Domain Knowledge and Meta-classification. J. Med. Biol. Eng. 33(4), 406–414 (2013)