Pervasive and Mobile Computing 6 (2010) 482–495

Contents lists available at ScienceDirect

Pervasive and Mobile Computing

journal homepage: www.elsevier.com/locate/pmc

Detecting activities from body-worn accelerometers via instance-based algorithms

Nicola Bicocchi, Marco Mamei *, Franco Zambonelli

Dipartimento di Scienze e Metodi dell'Ingegneria, University of Modena and Reggio, Emilia 42100, Italy

ARTICLE INFO

Article history:
Received 23 March 2009
Received in revised form 19 February 2010
Accepted 16 March 2010
Available online 27 March 2010

Keywords:
Context-aware computing
Body-worn sensors
Motion classification

ABSTRACT

The automatic and unobtrusive identification of user activities is one of the most challenging goals of context-aware computing. This paper discusses and experimentally evaluates instance-based algorithms to infer user activities on the basis of data acquired from body-worn accelerometer sensors. We show that instance-based algorithms can classify simple and specific activities with high accuracy. In addition, due to their low requirements, we show how they can be implemented on severely resource-constrained devices. Finally, we propose mechanisms to take advantage of the temporal dimension of the signal, and to identify novel activities at run time.

1. Introduction

The automatic and unobtrusive identification of user activities from sensor data is one of the most challenging goals of context-aware computing [1]. Several opportunities can arise from the availability of such information. For example, a service would be able to flexibly adapt its behavior to different circumstances (e.g., a smart phone application could turn silent, once recognizing that the user is in a theater watching a movie). As another example, the simple description of user activities could automatically produce entries in blogs, diaries and social network sites [2].

While advances in hardware technologies (e.g., smart phones, wireless sensors and RFID tags [3]) are making it feasible to collect a vast amount of information about the user in an unobtrusive way, it is still difficult to organize and aggregate all the collected information in a coherent, expressive and semantically rich representation. In other words, there is a gap between low-level sensor readings and their high-level context description [4].

In this paper we tackle the problem of turning low-level accelerometer data acquired from body-worn sensors into high-level activity descriptions. Body-worn accelerometers are suitable for this kind of application: they are small and economic, could be easily integrated in everyday objects and clothes, and they are already integrated in a number of smart phones.

To efficiently classify human activities on the basis of accelerometer data, we focus on instance-based mechanisms. Instance-based classifiers do not build an explicit, declarative representation of the class of interest but rely on the class labels attached to the training samples similar to the test sample to be classified. These methods have thus been called lazy learners, since "they defer the decision on how to generalize beyond the training data until each new query instance is encountered" [5]. We focused on these algorithms in that they are easy to implement, they offer state-of-the-art classification performance, the classification procedure can be easily understood (it does not rely on black-box models like in neural networks), they support on-line classification and training and, as detailed in the following sections, they can be implemented on resource-constrained devices.

Although the problem of classifying user activities on the basis of body-worn accelerometers has already been investigated in several other studies [6–9] (see Section 2 for details), the contribution of this paper is to show that simple

* Corresponding author.
E-mail address: mamei.marco@unimore.it (M. Mamei).

1574-1192/$ – see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.pmcj.2010.03.004

instance-based algorithms can classify accelerometer data more accurately than several other approaches employing more complex algorithms. In addition, we show that instance-based algorithms can be executed on resource-constrained devices like sensor motes.

The paper presents the following contributions and insights:

1. We show that simple instance-based (IB) algorithms offer state-of-the-art classification performance in this scenario (80% classification precision and recall with 16 classes). In addition, we verified that they can be easily implemented on resource-constrained sensors. In particular, the mainstream argument that an IB algorithm consumes too much memory because it stores individual feature vectors (see later) is unfounded. In this scenario, an IB algorithm can be implemented on memory-constrained devices keeping state-of-the-art performance.
2. We highlight that body-worn accelerometers are not only useful for recognizing simple motion/ambulation states (e.g., walking, running) like in most of the related work. These sensors can also recognize and discriminate among specific activities in other scenarios (e.g., use a PC, stir pasta, drive the car). This supports the use of accelerometer data and IB algorithms to fulfill the context-awareness vision.
3. We present a mechanism allowing IB classification to take into consideration the time series of the acceleration data and to aggregate classification results over a given time frame to considerably improve classification precision and recall.
4. We present a mechanism to introduce and identify new activity classes at run time. One of the main problems of IB classifiers is that they classify any new pattern as one of the predefined classes, no matter how different the new pattern is from those in the training set. Instead, practical systems require a mechanism to understand that a pattern deviating from those seen previously could represent a completely new activity rather than a different "gait" for the predefined ones.

Accordingly, the remainder of the paper is organized as follows. Section 2 surveys related work in this area. Section 3 presents the instance-based algorithms we used to classify activities using acceleration data and illustrates the possible applications enabled by these algorithms. Section 4 presents some novel mechanisms we introduced to better deal with the sensors' time series and to enable the on-line discovery of new classes. Section 5 describes the collection of the dataset. Section 6 presents the classification performances of our proposal. Section 7 details some results illustrating the suitability of the proposed mechanism to resource-constrained sensors. Section 8 concludes.

2. Related work

The problem of recognizing user activities on the basis of accelerometer data has been studied in a large number of works. Basically, all the approaches share the same basic framework:

1. A number of accelerometers are worn by the user typically at specific positions on her body.
2. Data is collected on a device (smart phone, laptop) and the user is asked to label his/her activities to provide ground-truth information.
3. Feature vectors are extracted from the raw accelerometer data.
4. A pattern recognition algorithm (whether instance based or using other mechanisms, such as the Hidden Markov Model and Support Vector Machine) is used to teach a model to recognize user activities on the basis of the feature vectors.
5. Once this training phase is concluded, new data is collected while the user performs daily activities. Such data is converted into feature vectors and then classified.
6. Some standard measures (e.g., accuracy, confusion matrices) are used to evaluate classification performances.

Since the basic framework of all the approaches in the area is almost identical, we can present their setting and results in a single table (see Fig. 1). For each of the related works being surveyed, we describe the number of accelerometers being used, where they have been located, the number and kind of activities the system tries to recognize, the classification mechanism being used, and the resulting classification accuracy.

Looking at the table, it is possible to see that, apart from a few proposals [10,6], most of the literature recognizes a rather limited set of ambulatory activities. It is also worth emphasizing that in [10] five accelerometers are used instead of three. This can obviously provide more data and ease the classification process. On the other hand, [6] exhibits large variance in classification precision and recall. Despite a high number of activities being considered, some of them are poorly recognized (less than 4% accuracy).

Although this table is useful for presenting several coherently related approaches, it does not allow for simply comparing the reported accuracies and deducing that one approach is better than another. These approaches are applied to different datasets, and thus it is not correct to compare their results: it can be more difficult to classify two similar activities than to separate 20 "disjunct" ones. Similarly, one approach can capture several *expressions* of the same activity (e.g., different ways of *running*), while another one might only capture only a few. Despite these limitations, the table gives an overall view of the classification task being addressed by the related work and of the accuracy obtained.

Our approach demonstrates that simple IB algorithms can reliably recognize a fairly large number of (also not ambulatory) activities.

We want also to emphasize that the use of IB approaches has recently gained a lot of attention in research and applications in several other areas.