Expert Systems with Applications 41 (2014) 6067–6074

Contents lists available at ScienceDirect

Expert Systems with Applications

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

Unsupervised learning for human activity recognition using smartphone sensors

Yongjin Kwon, Kyuchang Kang, Changseok Bae *

SW·Content Research Laboratory, Electronics and Telecommunications Research Institute, 218 Gajeong-ro, Yuseong-gu, Daejeon 305-700, Republic of Korea

ARTICLE INFO

Keywords:
Human activity recognition
Unsupervised learning
Healthcare services
Smartphone sensors
Sensor data analysis

ABSTRACT

To provide more sophisticated healthcare services, it is necessary to collect the precise information on a patient. One impressive area of study to obtain meaningful information is human activity recognition, which has proceeded through the use of supervised learning techniques in recent decades. Previous studies, however, have suffered from generating a training dataset and extending the number of activities to be recognized. In this paper, to find out a new approach that avoids these problems, we propose unsupervised learning methods for human activity recognition, with sensor data collected from smartphone sensors even when the number of activities is unknown. Experiment results show that the mixture of Gaussian exactly distinguishes those activities when the number of activities k is known, while hierarchical clustering or DBSCAN achieve above 90% accuracy by obtaining k based on Caliński–Harabasz index, or by choosing appropriate values for ε and *MinPts* when k is unknown. We believe that the results of our approach provide a way of automatically selecting an appropriate value of k at which the accuracy is maximized for activity recognition, without the generation of training datasets by hand.

1. Introduction

With an increasing interest in human health, it is necessary to obtain objective information of a patient to provide appropriate healthcare services. Although wearable sensors are convenient and useful for obtaining undistorted information from a human body, they may sometimes become an obstacle for healthcare services. Since most wearable sensors are attached directly to the user's body, some people may feel a sense of inconvenience. These disadvantages may generate incorrect results from healthcare services, making people refuse such services.

The inconvenience of wearable sensors can be resolved through the use of mobile devices (Anguita, Ghio, Oneto, Parra, & Reyes-Ortiz, 2012). Mobile devices such as smartphones and music players have recently become pervasive to the point that many people carry them at all times. Mobile devices usually incorporate various sensors, such as GPS sensors, accelerometers, or gyroscopes. For instance, most Android-powered smartphones have built-in sensors, and these sensors can be controlled by developers using Android APIs. Thus, instead of attaching a wearable sensor to the user's body, a mobile device can be kept in the pocket of the user's pants, which is less bothersome.

One of the most impressive studies on the sensors used in mobile devices is activity recognition. Human activity recognition is promising research that has been widely studied in recent decades. The sensors used in mobile devices can provide useful information for activity recognition. In particular, accelerometers, which are used as a source of fundamental information in many studies on activity recognition, are included in most smartphones (Brezmes, Gorricho, & Cotrina, 2009). If methods to recognize a user's activity with a mobile device can be achieved, it will be possible to develop many useful healthcare applications. For instance, we can monitor the activity states of a user, and can aggregate such activity states over time to obtain daily, weekly, and monthly ratios of activities. These ratios can be used to determine whether a user exercises regularly or sits for too long. Based on the estimated activity ratios, the application can recommend an appropriate activity to the user, such as walking outside or stretching. If we think beyond healthcare services, we may come up with more diverse useful applications.

Most studies on activity recognition have utilized accelerometers in either wearable sensors or smartphones. Interestingly, most such studies have considered activity recognition as a supervised learning problem. In fact, it is natural to think of activity recognition as a supervised learning problem because activity recognition

* Corresponding author. Tel.: +82 42 860 3816.
E-mail addresses: scocso@etri.re.kr (Y. Kwon), k2kang@etri.re.kr (K. Kang), csbae@etri.re.kr (C. Bae).

http://dx.doi.org/10.1016/j.eswa.2014.04.037
0957-4174/

classifies a given sensor dataset based on the activities. For the supervised learning, one of the most important things is the training dataset. The generation of a training dataset, however, is a tedious and labor-intensive work. Furthermore, the training dataset has some drawbacks. First the number of sensor records may be huge. For instance, when sensor data are recorded at a sampling rate of 50 Hz, the number of sensor records for an hour is 180,000. It is time consuming to label the whole records. Second it is difficult to remember the activities performed at a specific time. Especially, for short periods of an activity or at the boundary of consecutive activities, it is difficult to assign the correct activities. Last when the number of activities to be recognized varies, the training dataset should be regenerated. For these reasons, we need to seek new approaches of activity recognition without generating training dataset.

In this paper, we propose activity recognition using unsupervised learning assuming that the number of activities k is unknown. Although there are a few studies that have applied unsupervised learning approaches, they are inadequate to discuss the effectiveness of unsupervised learning for activity recognition, especially when k is unknown. Hence, we present experiments that examine different types of unsupervised learning algorithms to show that our approach can find an appropriate set of k at which the accuracy is maximized and can separate different activities. We first collected a series of sensor data from smartphones as the users performed five activities: walking, running, sitting, standing, and lying down. We then generated a list of feature vectors by aggregating the sensor data over sliding windows. To verify the usefulness of unsupervised learning techniques, we examined three clustering algorithms while assuming that the number of clusters is known, and observed whether they divided the vectors into five clusters precisely. We then investigated four clustering algorithms by assuming that the number of clusters is unknown to see whether they can still be applied to any series of sensor data, which are collected during an arbitrary number of activities. Hence, we observed whether unsupervised learning approaches can play an important role in activity recognition in future works.

The rest of the paper is organized as follows. In Section 2, we present the previous approaches on human activity recognition. Section 3 explains the details of the experimental setup, feature extraction, and descriptions of the sensor data for each activity. Section 4 shows the experiment results of the unsupervised learning techniques. Finally, Section 5 provides some concluding remarks regarding this research.

2. Related work

Many investigators have tried to recognize human activities using various combinations of sensors, which are included in cameras (Uddin, Thang, Kim, & Kim, 2011), wearable computers, and mobile devices. Accelerometers are common sensors for activity recognition because the accelerations measured rely on which activity the user performs (Mathie, Coster, Lovell, & Celler, 2004). Therefore, activity recognition has been studied using a number of accelerometers or with a combination of accelerometers and other types of sensors.

In some researches, attempts at using multiple accelerometers attached to different locations have been progressively conducted. The authors in Veltink, Bussmann, de Vries, Martens, and van Lummel (1996) performed a number of experiments that used two or three uniaxial accelerometers to distinguish several activities, including standing, sitting, lying down, walking, ascending stairs, descending stairs, and cycling. The researchers in Aminian et al. (1999) studied whether activities (lying down, sitting, standing, and walking) can be recognized using two accelerometers, one attached to the chest and the other to the rear of the thigh. In Foerster and Fahrenberg (2000), three uniaxial accelerometers were strapped to the sternum, and two uniaxial accelerometers were located on the left and right thighs to detect four basic activities (sitting, standing, lying down, and moving). Using only two accelerometers, the authors in Laerhoven and Cakmakci (2000) identified seven activities, sitting, standing, walking, running, climbing stairs, descending stairs, and riding a bicycle. The researchers in Bussmann et al. (2001) provided a technical description of an activity monitor in which four uniaxial accelerometers and one biaxial accelerometer were used to recognize activities such as standing, sitting, walking, climbing up, climbing down, cycling, driving, running, and laying down. The authors in Mantyjarvi, Himberg, and& Seppanen (2001) tried to recognize different moving activities, such as walking on a level surface, walking downstairs, walking upstairs, and not walking using two sets of accelerometers. In Bao and Intille (2004), sensor data were collected from 20 individuals wearing five biaxial accelerometers while doing twenty activities to show that a decision-tree classifier can recognize such activities with reasonable accuracy. The researchers in Krishnan, Colbry, Juillard, and Panchanathan (2008) examined five activities, sitting, standing, walking, running, and lying down, using two accelerometers. In Krishnan and Panchanathan (2008), the authors collected data from ten subjects wearing three accelerometers to identify seven activities, walking, sitting, standing, running, bicycling, lying down, and climbing stairs. In Mannini and Sabatini (2010), some experiments were performed that are similar to those in Bao and Intille (2004) in that they also used five biaxial accelerometers to collect the sensor data from 20 individuals, but applied classifiers based on Hidden Markov Models. The authors in Banos, Damas, Pomares, Prieto, and Rojas (2012) measured accelerations using a set of accelerometers placed on the hip, wrist, arm, ankle, and thigh to recognize four activities, walking, sitting, standing, and running. In Zhang, Liu, Zhu, and Zhu (2012), the authors varied the number of accelerometers, with different settings, and examined eight activities, standing, walking, running, jumping, lying, sitting, tooth brushing, and eating.

Some studies have combined accelerometers with other sensors, such as gyro sensors, microphones, and digital compasses. The researchers in Foerster, Smeja, and Fahrenberg (1999) tried to recognize nine activities, sitting, standing, lying, sitting and talking, sitting and operating a computer keyboard, walking, going up stairs, going down stairs, and cycling, using four accelerometers and some additional channels such as a microphone and an electrocardiogram. The authors in Lee and Masc (2002) created a system that uses a biaxial accelerometer, a gyroscope, and a digital compass to identify the user's location and activities, such as sitting, standing, walking on level ground, and going up and down a stairway. In Najafi et al. (2003), the authors utilized two accelerometers and one gyroscope on the chest to identify whether elderly persons were standing, sitting, walking, or lying down. In Parkka et al. (2006), the authors built a system that measures two accelerations using two accelerometers (one on the chest and the other on the wrist) and 16 different quantities with 20 additional sensors to recognize such activities as lying down, sitting, standing, walking, Nordic walking, running, rowing, and cycling. The authors in Subramanya, Raj, Bilmes, and Fox (2006) addressed similar activities by building a model using data from a triaxial accelerometer, two microphones, phototransistors, temperature and barometric pressure sensors, and GPS to distinguish between a stationary state, walking, jogging, driving a vehicle, and climbing up and down stairs. In Tapia et al. (2007), five accelerometers and a heart rate monitor were incorporated to automatically recognize activities with different intensities (lying down, standing, sitting, walking, running, etc.). The authors in Banos,