

Personalized Feedback based on Automatic Activity Recognition from Mixed-Source Raw Sensor Data

Harm op den Akker

Roessingh Research and Development
h.opdenakker@rrd.nl

Val Jones

Telemedicine Group
University of Twente
v.m.jones@utwente.nl

Hermie Hermens

Roessingh Research and Development
h.hermens@rrd.nl

Keywords: Wireless sensor networks, data mining, sensor data fusion, activity monitoring, COPD, biosignals.

Abstract

We present a data set consisting of multiple wireless sensors that monitor movement and various types of bio signals, recorded from patients that suffer from Chronic Obstructive Pulmonary Disease (COPD). From this data, the goal is to derive appropriate feedback to the patient that will motivate them to achieve a healthy lifestyle and a good physical condition.

1 Introduction

The AAL¹ project IS-ACTIVE (Inertial Sensing Systems for Advanced Chronic Condition Monitoring and Risk Prevention) started in April 2009. The project addresses continuous monitoring of activities and health status of patients, suffering from Chronic Obstructive Pulmonary Disease (COPD), in their daily environment. The goal is to promote a healthy lifestyle by providing personalized feedback on daily life activities taking into account the limitations for the patient caused by his chronic condition. To achieve this goal, we need to know what the patient is doing, and what the condition of the patient is throughout the day.

The patient will be equipped with a series of smart wirelessly networked sensor nodes. The final selection of sensors to be used has not yet been made but will likely include MEMS accelerometers, tilt switches, gyroscopes and magnetic compasses. Each sensor node will also include a microcontroller which takes care of sampling and networking, but resources must also be reserved for intelligent processing of the sampled data on the microcontroller itself. Besides the motion sensors needed for activity recognition, the patient will also wear biosensors to monitor his health status. Physiological parameters of interest include heart rate, some measure of respiration and oxygen saturation. In addition, analysis of audio recordings may be used to detect respiratory difficulty indicated by coughing and heavy breathing.

The resulting dataset will include sensor data outputs captured while performing a wide range of movements like walking, nordic walking (with sensors on the sticks), cycling, and any physiotherapy exercises that are commonly prescribed to COPD patients. Continuous series of sensor data will be hand-annotated with descriptions of the activities performed. If possible, the corpus will include video recordings of some basic activities, such as walking on a treadmill while wearing all the different sensors, so that these activities can be studied in greater detail afterwards. Because initially all raw sensor outputs are saved to the corpus, and the sampling frequency of the sensors will be set to a high level (100Hz or higher), the total quantity of data collected will be very large. The use of multiple movement sensors, such as the Inertia ProMove sensor², which will feature 9 degrees of freedom from three sensors: 3D accelerometer, 3D magnetic compass and 3D gyroscope, will result in a data set comprising as many as 40-50 layers of data.

The data mining challenge will be to automatically classify time segments of sensor data as belonging to one of the identified activities, and at the same time to calculate an estimate of the amount of strain that is put on the patient while performing that activity. The definition of strain in this context depends heavily on the individual patients, and can be seen as physical- or psychological strain, stress or a combination thereof. An important constraint on the algorithms to be designed is that they should run on the wireless sensor network nodes in a distributed way. Data transmission should be kept to a minimum to preserve battery lifetime, while processing power on the nodes themselves is limited. For examples of a distributed activity recognition approach see [Marin-Perianu *et al.*, 2008; Amft *et al.*, 2007].

The challenge in designing the automatic activity recognizers can be described by three requirements. First, the algorithms should use as little data as possible from the sensors in order to minimise the number of sensors actually needed and to enable reduction of each individual sensor's sampling frequency. Second, processing of individual sensor outputs should be done on the wireless sensor node itself, as far as possible, in order to reduce the need

¹Ambient Assisted Living

²<http://www.inertia-technology.com/>

for wireless transmission between sensor nodes. And third, the part of the algorithm which combines the various sensor outputs should be as simple as possible so as to be able to run on one of the (resource poor) nodes.

2 Approach

Because of the distributed nature of the task, we propose a layered approach with well defined subtasks that can be performed on specific nodes in the network. At the lowest layer, feature extraction from the wireless sensor data will take place. Once it is clear which features are needed for the activity recognition task, these features should be extracted on the sensor nodes themselves so that network transmission can be kept to a minimum. Then, the feedback device, which will most probably be some sort of PDA, is charged with the task of collecting relevant features from the nodes and doing the actual activity recognition. Note that with state-of-the-art PDA devices, the resources available for this part of the algorithm might not be that limited at all, but battery usage remains an issue. A similar approach is required for the biosensors, which will send their data (e.g. heartrate) to the feedback device on a previously defined minimum need basis. A second algorithm running on the feedback device will then combine biosensor and activity data and generate appropriate feedback for the patient.

This feedback is meant to help patients to be as active as possible, while preventing attacks of breathlessness. In order to provide each patient with the optimal feedback, the system will adapt to the behaviour and health status of the user. If, for example, a patient repeatedly chooses to ignore advice from the system to take it easy, with no serious health consequences for the patient, the system should become less cautious and allow the user to be more active. More importantly, if the system fails to warn a patient in time to lower his/her activity level, the system should issue its warnings more quickly. This general activity monitor is one of the envisioned applications, one that requires only a rather broad measure of activities. A second application is to aid patients in performing their daily physiotherapy exercises in a correct way. This requires a better accuracy from the activity recognition algorithm, because it has to correctly detect, for example, short series of arm or leg movements. The feedback device can then take on the task of personal coach, by keeping track of the exercise schedule while giving motivational feedback.

These different applications impose different requirements on the classification tasks. On the broad scope, the system should never mistake running for lying in bed, but mistaking slowly riding a bike for walking might not be a huge problem. On the other hand, the need for accuracy greatly increases when trying to detect all the actions that are performed in a physiotherapy exercise session. These differences have to be taken into account when collecting the training data for the algorithms. For detecting a bicycle ride, it may be sufficient to annotate a 15 minute trip from home to work as “riding a bike” (without indicating a stop for a red light, or the speed at every moment) and use it for

training. For the exercise patterns, it is probably a good idea to make video recordings of various sessions, and letting each phase of the movements be annotated by multiple annotators according to a previously agreed upon annotation manual. The inter annotator agreement then needs to be high overall, but small inconsistencies near the boundaries of movements may be acceptable.

For the annotation schema we propose a layered approach. On the highest layer, at least five different classes will be distinguished including riding a bike, walking, jogging, doing exercises and *non-active*. Then in the second layer, more detailed activities like the exact exercises can be annotated. These might include up to 10 different at home exercises for COPD patients. If necessary for the classification algorithms, a third layer may contain annotations of specific arm- or leg movements.

3 Feedback

As stated earlier, the final goal of the research is to promote a healthy lifestyle for COPD patients. We attempt to achieve this by providing feedback that motivates each individual patient to improve their physical condition to the maximum of their abilities. This raises the question of when and how to provide feedback, which is a non-trivial and not well understood issue. That is why an important part of our research will focus on using the recognized activity patterns and bio-signal data as input to a feedback system. This system can be seen as a sort of Clinical Decision Support System that will also have to adjust its ‘decisions’ (i.e. feedback responses) to how the patient reacts to them. At this point however, the details of the development of such a system are largely unclear.

To conclude, the goal of this article is to start a discussion on how to use data mining or machine learning techniques to eventually derive appropriate patient feedback from a large set of raw sensor data.

References

- [Amft *et al.*, 2007] Oliver Amft, Clemens Lombriser, Thomas Stiefmeier, and Gerhard Tröster. Recognition of user activity sequences using distributed event detection. In *EuroSSC 2007: Proceedings of the 2nd European Conference on Smart Sensing and Context*, volume 4793 of *Lecture Notes in Computer Science*, pages 126–141. Springer, October 2007.
- [Marin-Perianu *et al.*, 2008] Mihai Marin-Perianu, Clemens Lombriser, Oliver Amft, Paul Havinga, and Gerhard Tröster. Distributed activity recognition with fuzzy-enabled wireless sensor networks. In *Proceedings of the International Conference on Distributed Computing in Sensor Systems*, pages 296–313, 2008.