Pattern Recognition Letters 73 (2016) 33–40

Contents lists available at ScienceDirect

Pattern Recognition Letters

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

Optimising sampling rates for accelerometer-based human activity recognition☆

Aftab Khan[a,*], Nils Hammerla[b], Sebastian Mellor[b], Thomas Plötz[b]

[a] Telecommunications Research Laboratory, Toshiba Research Europe Limited, Bristol, UK
[b] Open Lab, School of Computing Science, Newcastle University, UK

ARTICLE INFO

Article history:
Received 3 July 2015
Available online 14 January 2016

Keywords:
Activity recognition
Accelerometers
Sampling rates
Statistics

ABSTRACT

Real-world deployments of accelerometer-based human activity recognition systems need to be carefully configured regarding the sampling rate used for measuring acceleration. Whilst a low sampling rate saves considerable energy, as well as transmission bandwidth and storage capacity, it is also prone to omitting relevant signal details that are of interest for contemporary analysis tasks. In this paper we present a pragmatic approach to optimising sampling rates of accelerometers that effectively tailors recognition systems to particular scenarios, thereby only relying on unlabelled sample data from the domain. Employing statistical tests we analyse the properties of accelerometer data and determine optimal sampling rates through similarity analysis. We demonstrate the effectiveness of our method in experiments on 5 benchmark datasets where we determine optimal sampling rates that are each substantially below those originally used whilst maintaining the accuracy of reference recognition systems.

1. Introduction

When Marc Weiser developed his vision of pervasive computing, one of the key promises of it was the prospect of disappearing technologies that "weave themselves into the fabric of everyday life until they are indistinguishable from it" [1]. The research fields of ubiquitous and pervasive, as well as wearable computing have since then matured rapidly and many of the original visions have become a reality for many – massively accelerated with the advent and ubiquitous uptake of smartphones and, more recently, wearable sensing and computing platforms such as smart watches. Smart environments, living labs, and especially mobile computing now constitute the central paradigm of what many call the third generation of computing [2]. As an enabling technology, automatic inference of the context and especially of the activities humans are engaged in – typically referred to as Human Activity Recognition (HAR) – plays a central role in the majority of ubiquitous and mobile computing applications.

Over the years a multitude of application scenarios have been developed for activity recognition – with the majority using wearable, miniaturised inertial measurement units (IMUs). Of particular interest have been tri-axial accelerometers that are used in such diverse domains like novel interaction techniques [3], situated support in smart environments [4], automated health assessments [5,6] or health care automation [7–9] to name but a few. Beyond the mere recognition of certain activities of interest, a number of application domains now even require the analysis of the quality of a person's activities [10]. Such skill assessment is, for example, relevant for progress assessment in physical rehabilitation [11], or for coaching in certain sports [12].

Real-world applications of human activity recognition require careful configuration of wearable accelerometers to balance accurate and reliable analysis of movement data with practical constraints surrounding battery life and storage requirements. The sampling rate, i.e., the temporal resolution at which tri-axial acceleration is measured represents the most critical parameter. It directly affects power consumption, data storage, and power or bandwidth requirements in case of wireless transmission. The dilemma lies in the fact that poorly chosen sampling rates can either lead to excessive energy and memory consumption, or, conversely, to missing detail of the activity data to be analysed. To avoid this, the decision of choosing an optimal sampling frequency is based on the balance between the information content against the amount of time a sensor could record for. For example, in various applications, accelerometry data is recorded for a long duration and it is crucial to be able to reliably determine the logging time whilst also considering the amount of information content.

In this paper, we investigate optimised sampling rates with a statistical approach that analyses the properties of raw

☆ This paper has been recommended for acceptance by R.Davies.

* Corresponding author. Tel.: +44 7817634757.
E-mail address: aftab.khan@toshiba-trel.com, aftab.khan@newcastle.ac.uk, aftabkhan00@gmail.com (A. Khan).

http://dx.doi.org/10.1016/j.patrec.2016.01.001

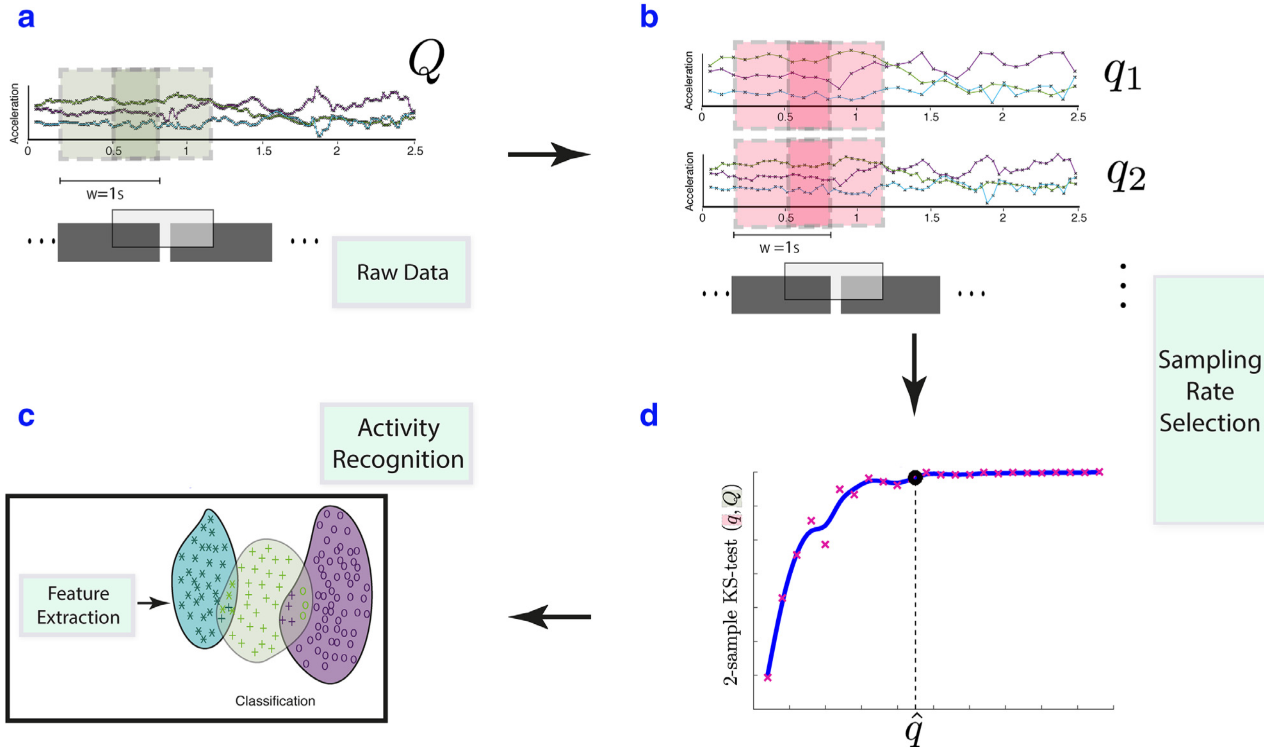

Fig. 1. Optimised sampling rate selection procedure presented in the context of an activity recognition system. Q represents the original sampling rate of the raw data (a); q_i represents lower sampling rates (b) while $\hat{q}$ is the optimal sampling rate calculated using a two-sample Kolmogorv–Smirnov (KS)-test between a dataset in its original sampling rate and subsampled versions of it (c). (d) represents the evaluation step of a standard activity recognition backend.

accelerometer signals (cf. Fig. 1 for an overview of the developed approach). We employ statistical tests for determining those sampling rates that are optimal for particular application scenarios. Through assessments on small, unlabelled sample sets sampling rates can be optimised automatically using our approach. In an experimental evaluation we demonstrate the effectiveness of our method on 5 different datasets, each covering a wide variety of human activities. We aim to successfully determine optimised sampling rates for each dataset, which is a pre-requisite for subsequent successful activity analysis. For a concrete evaluation, we also perform activity recognition on these datasets to show that the selected sampling rate is optimal in the given settings.

The developed method is universally applicable, i.e., without restrictions regarding the actual activities to be recognised in a deployment of human activity recognition systems. A possible application scenario would consist of running small scale pilot studies using the particularly chosen sensing solution to collect small amounts of – unlabelled – accelerometry data that is representative for the application domain. Through employing our analysis method practitioners could then automatically determine the sampling rate that is optimal for their application scenario and configure their sensing and analysis settings accordingly for most effective human activity recognition in their field. In our experience this resembles a typical deployment setting specifically in the health domain where non-experts often struggle to configure their sensing and analysis framework such that it produces optimal results.

2. Related work

Previous work suggests that sampling rates of approximately 20Hz are reasonable for “standard” human activities (e.g., [13,14]). Such standard activities usually correspond to periodic movements such as walking, running, or cycling, i.e., fairly regular and less complex movements. The motivation for sampling with said 20 Hz typically comes from the Shannon–Nyquist theorem [15], which states that for a successful, i.e., loss-less reconstruction of a particular signal, the data needs to be sampled with *at least* twice its highest frequency. It is assumed that voluntary human movements do not typically exceed 10 Hz and thus, according to Shannon–Nyquist, $\geq$ 20 Hz would be a reasonable choice when recording accelerometer data using wearable sensing platforms.

However, recent developments in HAR have moved the field towards the analysis of more complex movements and activities, and even beyond the recognition of activities [10]. For example, a number of methods now focus more on quality analysis of recorded activities and thus aim at a more fine-grained assessment of the details of the movement data. For such application cases a sampling rate that is too low would effectively hide essential details such that automatic signal analysis w.r.t. quality parameters and differences thereof becomes very difficult, if not impossible (e.g., mechanical vibrations may be insightful when assessing tool use [16]). In the case of human activities, $\geq$ 20 Hz *cannot* be simply considered as a natural sampling rate as it is difficult to reliably represent human activities. This could be linked with a lack of periodicity in accelerometry signals which also means that signals at this rate are significantly different to a much higher rate. It is for this reason that, in many cases, accelerometer data is now usually recorded at much higher sampling rates. For example, datasets like Opportunity and Daphnet Gait were recorded at 30 Hz [17,18], but in health and sports assessment scenarios 100 Hz was used [19,20], and in other domains sampling rates as high as 250 Hz have been used [21]. Certainly, the availability of inexpensive sensing hardware (such as *ADXL*335 tri-axial accelerometers) have also contributed to the wider use of higher sampling rates [22]. However, higher sampling rates come at a price for real-world deployments that rely on long term operation. The higher the sampling rate, the shorter the deployment time.