Available online at www.sciencedirect.com
ScienceDirect
Pervasive and Mobile Computing 3 (2007) 95–116
pervasive and mobile computing
www.elsevier.com/locate/pmc

# Recognition of emergent human behaviour in a smart home: A data mining approach

Sebastian Lühr*, Geoff West, Svetha Venkatesh

*Department of Computing, Curtin University of Technology, Kent Street, Bentley 6102, Western Australia, Australia*

Received 1 October 2005; received in revised form 2 August 2006; accepted 4 August 2006
Available online 11 September 2006

**Abstract**

Motivated by a growing need for intelligent housing to accommodate ageing populations, we propose a novel application of intertransaction association rule (IAR) mining to detect anomalous behaviour in smart home occupants. An efficient mining algorithm that avoids the candidate generation bottleneck limiting the application of current IAR mining algorithms on smart home data sets is detailed. An original visual interface for the exploration of new and changing behaviours distilled from discovered patterns using a new process for finding emergent rules is presented. Finally, we discuss our observations on the emergent behaviours detected in the homes of two real world subjects.


*Keywords:* Intertransaction association rules; Emergent behaviour; Visual data mining; Smart homes

## 1. Introduction

Current trends suggest that the global population will consist predominately of older people, those aged sixty and over, in as little as fifty years [1]. This demographic shift is expected to lead to an increase in the need for smart homes; intelligent environments

* Corresponding author. Tel.: +61 8 9266 7680; fax: +61 8 9266 2819.
*E-mail addresses:* luhrs@cs.curtin.edu.au (S. Lühr), geoff@cs.curtin.edu.au (G. West), svetha@cs.curtin.edu.au (S. Venkatesh).

1574-1192/$ - see front matter 
doi:10.1016/j.pmcj.2006.08.002

that are able to assist their occupants in maintaining independent lifestyles for as long as possible [2]. In our work, we seek to facilitate the creation of homes that are able to detect the presence of new, possibly abnormal, behaviour in their occupants and to take action accordingly. Appropriate action may be to query the occupant on the new behaviour, to jog their memory on a task that they were carrying out or even to alert an occupant's relative that their assistance is required. Detecting such deviations from normal behaviour requires models of a person's expected behaviour to compare with incoming data.

One of the more popular approaches favoured by researchers for modelling human behaviour has been the application of graphical models. Training and inferencing using these models is, however, computationally expensive and generally limited to applications in which it is reasonable to assume that human activity can be represented as sequences of asynchronous activities or events. The task of modelling human behaviour as precise sequences of events is made difficult, however, by our tendency to interleave our activities and to adjust our behaviour when we are interrupted. We propose a novel application of Intertransaction Association Rule (IAR) mining [3] as a means of tackling this issue. IARs are implication rules that allow us to capture the associative, non-sequential, relationship of events observed within a home while retaining some of the higher level temporal context in which these events occur.

The research contributions of this work are threefold. The first addresses the issue of mining the IARs present in the sensor event logs from a home. Such environments generate data sets that contain frequent occurrences of a large number of events over relatively short periods. This poses a problem for the current EH-Apriori [3] and FITI [4] algorithms for IAR mining as they rely on a computationally costly candidate-generation-then-test approach for rule discovery. This technique requires that $k$ passes over a database are made to retrieve the set of frequent rules up to length $k$. Each pass over the data requires the generation of candidates — the set of all *possibly* frequent associations given those found in a previous pass. The scalability of such algorithms is hence limited due to the computational complexity of generating and testing the frequency of a combinatorial number of candidates. The number of candidates generated at each pass $k$ of a worst case scenario given $n$ database items and an intertransaction window of length $w$ is $\sum_{r=0}^{k-1}\left[\binom{n}{k}\binom{nw}{r}\right]$. We propose the application of our Extended Frequent Pattern Tree (EFP-Tree) [5] as a means of tackling this issue. The EFP-Tree is an extension of the Frequent Pattern Tree (FP-Tree) [6] for IAR mining. It uses a divide and conquer approach to avoid candidate generation and requires only three passes over a database.

The second issue we face is how to gain insight into a person's behaviour so as to detect abnormality from an overwhelming number of rules that the mining process is likely to uncover. We introduce the use of emergent IARs as a novel means of finding patterns of behaviour that are of interest to us. Emergent IARs are those rules that display significant growth from one data set to another. Their presence may indicate abnormality evident as either a previously unseen pattern of events or unusually frequent occurrences of behaviour that would otherwise be considered normal. Emergent IARs offer a convenient means of identifying changes that would otherwise be difficult to discern through manual inspection of the rule sets. For example, the real world event logs from a single week of data can produce around 7200 patterns which can be distilled down to approximately 150 emergent rules that are likely to be of interest to us. The same data set mined with a