Expert Systems With Applications 43 (2016) 15–22

Contents lists available at ScienceDirect

Expert Systems With Applications

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

Training with synthesised data for disaggregated event classification at the water meter

M. Wonders [a,b,*], Z. Ghassemlooy [a], M. Alamgir Hossain [b]

[a] Optical Communications Research Group, United Kingdom
[b] Computational Intelligence Group, Faculty of Engineering and Environment, Northumbria University, Northumbria, NE1 8ST, United Kingdom

ARTICLE INFO

Keywords:
Training data synthesis
Machine learning
Load disaggregation
Assisted living

ABSTRACT

Activity recognition in monitored environments where the occupants are elderly or disabled is currently a popular research topic, with current systems implementing ubiquitous sensing or video surveillance techniques. Using disaggregated data from smart meters could be a viable alternative to what is often perceived as intrusive recognition technology. Disaggregation methods have proven to perform exceptionally well when trained with large quantities of data, but gathering and labelling this data is, in itself, an intrusive process that requires significant effort and could compromise the practicality of such promising systems. Here we show that by synthesising labelled training data, using a domain specific algorithm, an innovative water meter disaggregation system that uses Artificial Neural Networks (ANN), Support Vector Machine (SVM) and K-Nearest Neighbour (KNN) classifiers can be trained in minutes rather than hours. We show that by artificially synthesising labelled data accuracies of 83%, 79% and 85% with the SVM, ANN and KNN classifiers, respectively can be achieved. Though these values are marginally lower than 89%, 83% and 89% achieved with no synthesis, the measure of accuracy masks the underlying imbalance of representative classes in the data set.

1. Introduction

According to a report commissioned by Ofcom in 2010 (Lewin, Adshead, Glennon, Williamson, & Moore, 2010), there are around 1.8 million people with moderate to severe disabilities in the UK, with 72% of these aged 65 or above. The same report predicts that the requirements for health care within the home will increase by 16% over the next 20 years. The 2010 key issues for parliament report (Mellows-Facer, 2010) states that from an estimated UK population of around 61 million, 11.5 million people are of pensionable age and 4 million are living alone. It also states that 32 of these die alone and unnoticed each year. As the population ages these numbers are expected to increase and will undoubtedly put a substantial strain on care providers, whether formal i.e. National Health Service (NHS) or informal i.e. family, friends or neighbours. According to a 2009 Help the Aged report (Help the Aged, 2009) an increase in divorce rates, as well as smaller family units and the increased global mobility in family members, will make it much more difficult to maintain the current input from those involved in informal health care.

Telecare at home is proposed as one of the strategies that can help to minimise the impact as well as maintain the independence of an ageing population, and by incorporating more and more digital telecare and smart home technologies into living environments, there exists the potential to provide more flexible assisted living spaces for people with the special needs that come with age or disability.

Many of the studies carried out in this area, for example Lee and Mase (2002), Philipose, Fishkin, Perkowitz, Patterson, Fox, Kautz, and Hahnel (2004), Steinhauer, Chua, Guesgan, and Marsland (2010), Logan, Healey, Philipose, Tapia, and Intille (2007) and Tapia, Intille, and Larson (2004) focus around the use of ubiquitous sensitised spaces or video based surveillance and use activity recognition systems to interpret the data. These methods are generally intrusive and require expensive computational input as well as a high level of deployment impracticality.

We propose a system of non-intrusive sensor fusion that will implement disaggregation techniques, at both the water and electricity meters (Gupta, Reynolds, & Patel, 2010), to infer the activities of daily living (ADL) of the inhabitants. The data from each meter will be complemented with contextual data from simple passive infrared (PIR) sensors positioned in each room to form the system of sensor fusion. This contextual data from the PIR sensors is used to overcome the problem of similar signatures created by similar devices, by indicating the location of the device in use, for example, in terms of water

* Corresponding author. Tel.: +4401912273602.
E-mail addresses: martin.wonders@northumbria.ac.uk (M. Wonders), z.ghassemlooy@northumbria.ac.uk (Z. Ghassemlooy), almgir.hossain@northumbria.ac.uk (M. Alamgir Hossain).

http://dx.doi.org/10.1016/j.eswa.2015.08.033
0957-4174/

meter disaggregation, two similar toilets will have similar water use signatures but will very rarely be in the same room.

A limitation of the many proposed systems of water disaggregation such as those investigated by Chen, Kam, Zhang, Liu, and Shue (2005a), Fogarty, Au, and Hudson (2006), Froehlich, Larson, Campbell, Haggerty, Fogarty, and Patel (2009), Srinivasan, Stankovic, and Whitehouse (2011), Larson, Froehlich, Campbell, Haggerty, Atlas, Fogarty, and Patel (2012), Nguyen, Zhang, and Stewart (2013b), Nguyen, Stewart, and Zhang (2013a) and Nguyen, Stewart, and Zhang (2014) has been the lack of labelled training data available at both design and implementation time. Here we propose a method of training data synthesis, comprising of a small amount of labelled data, representing the extremes of fixture usage combined with artificially synthesised data, used to simulate the many possibilities between these extremes. This particular domain serves as a suitable test case for labelled data simulation as the extremes of data flow are limited by the practical application of the fixtures, for example a standard WC has a flush mechanism that is limited by the maximum capacity of the cistern, and likewise a wash basin or bath have a limited maximum capacity and useful daily functionality can be assumed with a fair degree of accuracy. The process of training data synthesis, with a view to maximising the quantity of training data, has been used in various studies of which Varga and Bunke (2003b) and Schuller and Burkhardt (2010) are examples and have shown to improve the performance in classification when there is a scarcity of labelled training data and hence deployment impracticality.

This paper documents the results of testing three different supervised machine learning techniques on water flow data gathered from a real environment, using a prototype water meter with a simple Hall Effect water flow sensor. This data is used to train and test ANN, SVM and KNN classifiers to determine the performance of each method using real data. We then implement a method of training data synthesis to produce artificial training data and proceed to train each classifier with this data. The models trained with synthesised data are then tested using the original real data and the performance of each model is compared with those using no synthesised training data.

2. Related work

There are several areas of related work with respect to the proposed system of activity recognition, here we will discuss only those areas relevant to water usage monitoring, as an activity recognition sensor.

Chen, Hsu, and Shen (2005b) carried out a study on automated bathroom activity monitoring using sound signals from microphones fitted in a small shower room. Hidden Markov Models (HMM) were used as the classification technique and produced an average classification accuracy of 83.55% over six activity types. There is an acknowledged issue of privacy, whether actual or perceived, in this work but the results are indeed encouraging. To produce the models used in the real trials, the authors used a training period of 10 days, representing a major practical limitation of the system. Fogarty et al. (2006) extended the idea of using microphones for activity recognition beyond the bathroom by fitting devices to the cold and hot water inlet as well as the main soil stack. This study reports an average recognition accuracy of 86.75% over 8 water usage event types. This technique alleviates the privacy issues highlighted in Chen et al. (2005b) by taking the microphones out of the private areas and into the basement (the location of inlets and outlets), but still suffers from the need for extensive labelled training data to create the classification models. This study also acknowledges a reliability issue when classifying events generated from similar fixtures, as well as issues relating to ambient noise pick-up on the microphones.

Froehlich et al. (2009) proposed a system that replaces the microphones of Chen et al. (2005b) and Fogarty et al. (2006) with a pressure sensor used to classify the pressure changes associated with the opening and closing of valves as individual fixture events. This method resolves the issues of privacy by replacing microphones with pressure sensors, and is also less intrusive as the sensors can be fitted to external water taps. The study reports extremely impressive results of 97.9% using a template classification model combined with a KNN classifier. The report acknowledges the deficiencies in classification of compound events but does go some way to classifying between same fixture types in different rooms. However, the work does not indicate the training data requirements of the system. The study was followed up by Larson et al. (2012) with an extended analysis of the technique as well as a comparison of the template model against an HMM classifier. This follow up study reports that both techniques produce aggregate classification accuracy of greater than 90%, but confirms that significant calibration on the training data is necessary and there still remain some limitations in the classification of compound events.

Srinivasan et al. (2011) solved the problem of classifying similar events in different locations by using data from PIR sensors to determine context. The study focus was primarily on water disaggregation to determine fixture level consumption, and reported 86% recognition accuracy over 467 fixture events. The work used a combination of Canny edge detection and unsupervised Bayesian inference and reported that no training data is required, but the fixture level identification process requires very specific flow rates and event durations and thus the need for extensive calibration.

Nguyen et al. (2013b) propose the development of an automatic flow trace analysis system and present a classification model that uses a Dynamic Time Warping (DTW) algorithm and Hidden Markov Models (HMM) along with event probability data to identify single fixture level events. The study uses labelled data gathered from 252 residential households over a period of 2 years during the summer and winter months. The published paper reports an average classification accuracy of 84.1% and shows potential when compared to a commercial package that takes the water efficiency and flow rate statistics of each fixture to determine usage. This study was followed up by Nguyen et al. (2013a) where the issue of classifying combined fixture level events is presented and tackled by incorporating gradient vector filtering to separate combined fixture level events into single events before using the HMM models for event classification. The study reports that approximately 88% of combined fixture level events were classified accurately, and also suggests that future research would require the collection and manual labelling of a much larger sample of combined events. The most recent publication by Nguyen et al. (2014) presents an adaptive system that can use the initial classification model learned from the 252 residential households and adapts this for new previously unseen households. The basic principle of the system is to use the current model to determine any fixture level events that cannot be classified, these events are then classified and added to the model and used to improve system performance. The study also discussed the development of a software package which, as presented, allows the occupants to manually modify the labelled data and upload this to the database to improve classification accuracy. Using this system the study reports that most of the achieved recognition accuracies for all end use categories were approximately 90%.

Like many other machine learning applications, all of the above methods of disaggregation require many samples of labelled training data to perform efficiently, but the practicalities of collecting this labelled data in a domestic environment are restricted by both time and access constraints. These restrictions on labelled data collection time and access can also result in labelled data that is unbalanced in favour of more frequent training examples and hence result in biased machine learning models. Synthesis of training data for machine learning applications suffering from a shortage of initial training examples has been used in emotion recognition (Schuller & Burkhardt, 2010) techniques, where speech is synthesised to train acoustic models for emotional speech recognition. Schuller et al. compared models