Understanding Health Changes Through the Analysis of Electricity Consumption Data

By:

Jennifer Salter

A thesis submitted in partial fulfilment of the requirements for the degree of Doctor of Philosophy

The University of Sheffield
Faculty of Social Science
Information School

October 2015

Abstract

With improvements in living standards and innovations in medical care, life expectancy has increased. However, although people are living longer, particularly in developed countries, they are not necessarily healthier during the additional years of life, with a rising number of people with long-term physical and mental health conditions that require supported living, for example, within a care home or hospital environment. In response to the rising economic costs of managing long term conditions, successive governments have developed policies to reduce the use of institutional environments (e.g., care homes) and of unplanned hospital admissions, and are encouraging the development of systems which aim to monitor, support and manage people's health in their own home.

These developments have lead to increased research on using remotely monitored, sensor-based technologies to provide relatives, carers and health care professionals with timely data about the well-being of older people living independently, and so provide timely and appropriate support effectively, thus helping them remain in their own homes, especially when they have long-term health problems.

The aim of the research described in this thesis was to investigate the use of an electricity monitor to recognise and monitor changes in resident's daily activities. This was achieved using two phases; the first conducted a survey to gather information about which activities and features that carers and relatives would like to have access to, so as to be reassured about their relative's health and well being. The second phase collected and analysed electricity consumption data from four households for a one-week period, to develop models to identify when specific activities had been undertaken, e.g., using the shower, using a kettle.

This research concluded that the monitoring of general and some specific activities is important to the relatives and carers, although the best form of reassurance about their relative's situation was felt to be human contact. Following the analysis of the electricity consumption data, it was concluded that while it is possible to recognise appliance usage from whole house electricity consumption data, the variability and lack of transferability between houses and appliances would mean that the large-scale use of this type of monitoring would require considerable further development.

Acknowledgments

I would like to extend my deepest gratitude to all those who have contributed and supported me through this research. With particular thanks and gratitude to:

My family and friends who have been my constant support through the years.

My supervisor Professor Peter Bath for his constant support, guidance and expertise that made the completion of this research possible.

My initial supervisors Dr Simon Tucker and Dr Simon Brownsell for their guidance through the first year of my PhD.

My friends and colleagues at the University of Sheffield Information School and Computer Science departments for their support and guidance. With a special thanks to Dr Stuart Wrigley for his assistance with the equipment set up.

And finally to all those who agreed to participate and made this research possible.

Table of Contents

List of Figures

List of Tables

List of Abbreviations

ACHE	Adaptive Control of Home Environment
ADL	Activities of daily living
AVG	Average
DVD	Digital versatile disk
ECG	Electrocardiogram
FN	False negative
FP	False positive
GPS	Global positioning system
GUI	Graphical user interface
HIV	Human immunodeficiency virus
ID	Identification
iDorm	Intelligent dormitory
IR	Infar-red
KDD	Knowledge discovery in database
LED	Light-emitting diode
NHS	National health service
NPV	Negative predicted value
PC	Personal computer
PPV	Positive predicated value
RECAP	Recognition of electrical appliance and profiling in real-time
RFID	Radio identification systems
RMS	Root mean squared
ROC	Receiver operating characteristic
SVM	Support vector machines
STD	Standard deviation
TN	True negative
TP	True positive
TV	Television

Chapter 1: Introduction

1.1. Introduction

With the increase in the age of the population and the growing economic cost of providing care for older people or those with long-term care and social needs, there has been increasing interest in looking at ways to support people to continue to live independently. One area of research has been to investigate ways of combining sensor and communication technologies to provide remote monitoring for older people or those with long-term health conditions, in their own homes. This thesis will present the research undertaken with the aim of investigating the use of an electricity monitor as a potential way of remotely monitoring the activities of an older person or those with long-term health conditions. This chapter has been split into a number of sections, with the background to this research described in sections 1.2-1.4. Section 1.5 will highlight the motivations behind conducting this research. Section 1.6 will provide the aims and objectives for the research, with section 1.7 showing the research questions. Finally, section 1.8 will provide an overview of the organisation of this thesis.

1.2. The ageing population

The United Kingdom, like many other developed countries, is facing an increasing ageing population. In 2010, one in six of the population of the United Kingdom were aged 65 or over and it is predicted that this will increase to one in four by 2050 (Cracknell, 2010). Globally, The United Nations (United Nations, Department Of Economic And Social Affairs, Population Division, 2013) estimates that in 2013 there were 841 million people aged 60 or over out of a population of 7.2 billion. This is predicted to increase to 2 billion out of a population of 9.6 billion in 2050 and almost 3 billion of a worldwide population of 10.9 billion by 2100. By 2100 the world's population of people aged 60 or over is estimated to have tripled.

The increase in age expectancy over recent years has brought many advantages to individuals as well as to society in general. The advantages, on a personal level, mean people are spending more time with their family and friends. Economically, people are working for longer meaning an increase in tax revenue and world output and companies are retaining their older workers again adding to companies' skills and resource and supporting growth.

Growth in population through increase life expectancy has benefits to society but there is also a significant social and economic cost of an aging population. Some countries where there is already a significant growth in the older population are starting to tackle some of the issues in both increased health and social cost. An example of this cost in the United Kingdom in 2010 was that 65% of the Department of Work and Pensions benefit expenditure went to those over the working age, equivalent to £100 billion or 7% of total public expenditure (Cracknell, 2010). There is also a health cost as well, with more money spent on retired households than on non-retired households by the National Health Service (NHS) (Cracknell, 2010).

An increase in population does not necessarily mean that people are healthier as they age. Older people with long-term health conditions are more likely to need care and also the more expensive that care will be to provide (Botsis & Hartvigsen, 2008). The Department of Health (Cracknell, 2010) estimated that it is three times more expensive to provide health and social care for the 85+ age group than for the 65-74 year age group. With the predicted increases in the population there will inevitably be an increase in the proportion of the population suffering from long-term health conditions or disabilities, for example, dementia, Type II Diabetes, cardiovascular problems etc.. These people will have long-term health and social care needs.

1.3. Smart homes, health smart homes and telecare

The increase in the ageing population, the rise in social and health costs and the increased prevalence of long term health conditions or disabilities have lead to an increase in research outlining the use of technologies to support health and social care. Examples of some these research areas are smart homes, health smart homes, telecare and the related disciplines. Within these areas, smart homes (as defined in section 2.3.2) and health smart homes (as defined in section 2.3.3) have described the development of homes that can be used to provide monitoring and support for their residents. Telecare and its related disciplines is developing technologies that can be placed or installed into people's homes to provide monitoring and support for the resident. The research within these areas can simplistically be split into two broad groups. The first area discusses providing monitoring and support for specific health conditions or disabilities - an example of work in this area is the work of (Lotfi, Langensiepen, Mahmoud, & Akhlaghinia, 2011) that discussed the use of smart home technologies to identify and predict abnormal behaviours in people with dementia and the second area outlines the providing of general monitoring and support for those in sheltered accommodation (Glascock & Kutzik, 2007).

Within the areas of smart homes, health smart homes, telecare and its related disciplines the use of different types of sensors combined with communication technologies such as the Internet (Glascock & Kutzik, 2007) or phone lines (Sixsmith, 2000) have been prominent areas of research. Sensors can be used to provide information and feedback about a large number of measurements and actions, for example, feedback about the resident's environments (temperature within the home (Intille et al., 2006), physiological data (e.g., heart rate (Agoulmine, Deen, Lee, & Meyyappan, 2011)) or physical activities (e.g., movement around their home (Helal et al., 2005))). Sensors can also take a number of different forms, for example, wearable sensors or fixed sensors attached to walls. The use of sensors allows many aspects of the resident's activities to be monitored and, combined with communication technologies, can provide feedback on activities to the resident's caregiver or relative.

The sensors used within this area of research can be split into two groups, active and passive sensors. Active sensors are those that provide an instant response to an emergency or change of situation, for example, a fall sensor that alerts a carer or emergency services that the resident has fallen over and cannot get back on their feet. In contrast, passive sensors are those that continually monitor the resident's activities or behaviour and the data are interpreted by using computer algorithms to identify potential changes that can be attributed to other factors, such as deteriorations in health over time. This then allows carers to act before more serious consequences arise.

The research within the area of telecare and its related disciplines has highlighted the benefits of using this technology for supporting elderly people, and those who have long-term health problems, to live independently in their own homes, rather than being hospitalised or institutionalised. Appropriate technology offers a more cost-effective system of caring for older people and those who are ill and reduces expensive hospital admissions (Sixsmith, 2000).

Although these areas of research have highlighted benefits, there are also some areas of concerns, such as privacy and the intrusiveness of the sensing technologies. Placing or installing sensor technologies into people's homes could be seen as intrusive to the resident, especially if it involves the use of cameras and microphones (Sixsmith et al., 2007) and also remind them of their vulnerabilities (Stowe & Harding, 2010). The privacy of the resident is also an issue, as the work of Perry, Beyer, & Holm, (2009) highlighted: although the intentions for monitoring are different, the principles used to monitor people within their own homes are the same as those used by security agencies and governments to monitor criminals or people of interest.

1.4. Electricity

Electricity is a commodity, which, until recently, the consumer was not generally aware of how much electricity they were using with specific appliances or in particular situations. Over recent years, several factors have led to an increased awareness of residential electricity usage. Examples of some of these are the increased cost of energy (Chetty, Tran, & Grinter, 2008), growing awareness about sustainability (Fischer, 2008) and government goals to cut CO_2 emissions (Climate Change Act 2008).

A recent study by the Department of Energy and Climate Change, the Energy Saving Trust and Department of Environment, Food and Rural Affairs has shown that, on average, a household wastes between £50 and £86 per year on appliances that are on standby or in a non-active state (Energy Saving Trust, 2012). One way to reduce fuel wastage is to provide the user with feedback into how much energy they are using, and on which appliances. Suggestions for in-home energy monitor displays to provide feedback to the user have been proposed and developed since the 1970s (Winett, Neale, & Grier, 1979). Recently, as described in the review by Kulkarni, Welch, & Harnett, (2011), a large range of different electricity monitors have been developed, which can be placed into the home, to provide real time feedback to the resident about their current energy usage in an effort to support the reduction of electricity consumption and costs.

The development of a wide range of electricity monitors that can be used to monitor a large number of variables about a person's electricity usage, combined with the ease and possibility of storing and transmitting the data over the Internet, have all opened up other opportunities beyond the goals of saving energy and providing the user with the knowledge about their energy usage. It is the potential use of these devices that is the focus of the research in this thesis.

1.5. The research and motivation

The research described in this thesis outlines the use of combining an electricity monitor with communication technologies, in this case the use of the Internet, to add to the current research being undertaken in the use of technologies to monitor people's specific activities as well as behaviour or lifestyle monitoring (Brownsell, Blackburn, & Hawley, 2008).

An electricity monitor, as described in section 1.4, can provide the user with an overview of their electricity use either currently or historically. An electricity monitor also provides granularity in its recording and, with analysis of the data, can potentially be used to show when different electrical appliances have been turned on and off. The data may also then be used to infer when undertaking a particular activity and, with longitudinal data collection, to highlight changes in activities.

The growing ease and availability of electricity monitors in recent years, as well as the relative simplicity of recording, storing, transmitting and analysing the data makes this a feasible area of research. The advantages of using an electricity monitor are that they are relatively cheap to buy, as they are widely available on the commercial market (as discussed in more detail in section 2.4) and are also easy to install. They also only require three pieces of equipment to be installed into a person's house and do not require home modifications or the installing of large amounts of sensors. The research has universal applications as almost all households use electrical appliances for some, or all, daily living tasks.

The ethical and privacy issues around using technology to support people in their own homes has surrounded the growth in research into using sensor technologies for support and monitoring. The use of an electricity monitor as an extension to the research already carried out in the area could provide a method of monitoring that can be perceived as less intrusive, or even non-intrusive, as the installation of an electricity monitor into a home requires no modification to the house. Sections 2.3.8 and 2.4.6 of this thesis provide a discussion into the ethical and privacy issues around placing monitoring equipment into a home but also the different views of those who have monitoring equipment placed in their homes. As well as the ethical and privacy considerations with the placing of monitoring equipment into a home, there is also a security consideration. The security of data, that has been recorded using sensors, is a growing area of concern and will be discussed in more detail in section 2.3.9.

As the evidence highlighted in section 1.2, with people living longer but not necessarily living healthier, there is an increase in the prevalence of chronic, i.e., long-term conditions. Whilst, for some chronic diseases, with good management little physical changes may occur over time, e.g., diabetes, arthritis, other diseases, such as dementia, are progressive diseases and will lead to deterioration in the patient's condition over time. The aims of many monitoring and support systems are to identify changes by using effective detection, so reducing the need for unplanned hospital admissions. Technology, in the form of telecare, has been developed to allow remote monitoring and care of individual patients; however, this is often only instigated after the first, or even second, unplanned admission has occurred and is often focused on detecting sudden, not chronic (i.e., long-term), changes. In addition, telecare systems may be felt to be intrusive, requiring people to actively wear sensors or regularly undertake potentially invasive tests. The use of monitoring of household electrical usage is both less intrusive and could also provide a low cost method to identify changes to a resident's activities, which would allow a more timely intervention.

1.6. Aims and objectives of the research

The overall aim of this project is to examine the potential use of electricity monitoring devices for monitoring the activities of older people or those with long-term health problems. The aim is to investigate if an electricity monitor could be used to monitor specific activities that may then be able to identify overall lifestyle or behavioural changes of the resident. More specifically, the objectives of this project are:

- To examine the views of relatives and carers about the use of sensors in the home and, more specifically, activities or tasks that the relative or carer believe are most relevant to be monitored.
- To examine the feasibility of collecting data from a single electricity monitor from multiple homes.
- To analyse the electricity data, from a number of households, to try to determine when different appliances have been used and hence infer different activities have been performed.
- To make recommendations for the use of a single electricity monitor as a remote monitor used to monitor specific activities as well as lifestyle of behavioural changes.

1.7. Research questions

For this research, the overall research question is, can measuring of electricity consumption data, whilst taking into account the needs for privacy, be effective in the monitoring of activities of older or chronically ill people?

This research question has been divided into a number of smaller research questions for the two parts of this thesis. The research questions relating to the survey are shown in section 1.7.1 and the research questions for the analysis of the electricity data are shown in section 1.7.2.

1.7.1. Survey research questions

As is discussed in the literature review (Chapter 2), there is limited research into the views of carers and relatives into the type of tasks or activities they would want a monitoring system to monitor and provide feedback on. It was decided to conduct a survey with the following research questions:

- What are the priorities of relatives and carers to the importance of knowing certain specific tasks and types of activities have been performed.
- What are the views of relatives or carers on how intrusive a remote monitoring system should be?
- What are the views of relatives or carers into, what should the properties of remote monitoring system be?
- Who, in the view of the carers/relatives should have access to information provided from a remote monitoring system?

The survey carried out to answer these research questions is described in Chapter 4.

1.7.2. Electricity analysis research questions

As is also highlighted from the literature review (Chapter 2), there has been limited research into collecting whole house electricity consumption data and analysing this data to determine when different appliances have been used. There has also been limited research into the transferability of the models developed to recognise appliances usage as well as analysis of the differences between whole house electricity consumption usage across multiple houses. The research questions for the analysis of the whole house electricity consumption data are:

- What is the feasibility of collecting whole house electricity consumption data from multiple homes?
- Is it possible to accurately recognise appliance usage from a single whole house electricity consumption data?
- How transferable is the developed model across multiple households' electricity consumption data?
- What are the differences, if any, between appliances and households from the multiple whole house electricity consumption data?

The study carried out to answer these research questions is described in Chapters 5 and 6.

1.8. Organisation of the thesis

This thesis is organised into a number of chapters.

Chapter 2 of this thesis provides the literature review and is divided into two sections. The first of these sections describes the literature into home monitoring and, more specifically, smart homes, health smart homes and telecare. This section also gives an overview of the large number of different sensors that have been used by researchers for monitoring purposes. In addition, it discusses a number of ethical and privacy considerations associated with the use of monitoring technologies to monitor resident's activities within their homes. The second section describes current research into the use of electricity monitors for appliance and activity recognition. This section also gives an overview of the different types of electricity monitors available, as well as ethical and privacy considerations with their use.

Chapter 3 describes the methodology for this study and used in this research. This chapter gives a detailed description of the methods used for the data collection and analysis of the survey data collected as the first part of this thesis. This chapter also provides a detailed description of the collection and analysis of the electricity consumption data and appliance diary data, which was collected as part of the second half of the study described in this thesis.

Chapter 4 describes the survey undertaken for the study; it describes the analysis of the data collected from the survey as well as the methods used to analysis both the qualitative and quantitative sections of the survey. This chapter also presents the results from this survey and discusses the results in relation to the development of a model to recognise activities from electricity usage.

Chapter 5 describes the trial analysis of the electricity data; it describes the process used to collect the electricity and associated diary data and how they were pre-processed into a format that could be used for analysis. The chapter then describes a number of different trials used to develop a method to recognise appliance usage from the electricity consumption data from one house. It then presents the final method adopted, as well as discusses some of the issues and limitations found during the analysis of the data from the first house.

Chapter 6 provides the results and discussion of the analysis of the electricity consumption data from three further houses, using the method developed in Chapter 5. This chapter also presents a discussion of the results from all four houses, as well as the issues and limitations with the analysis of electricity consumption data.

Chapter 7 provides the conclusion of this thesis. This chapter summarises the main findings from this research and how they relate to the literature as well as the original aims, objectives and research questions. This chapter also offers an overview of the limitations of this work and finally discusses the areas of further work.

1.9. Conclusion

This chapter has given a brief overview of the context of this research and the motivations behind this area of work. This chapter has also presented the aims and objectives of this research as well as the structure and the layout of this thesis. The next chapter of this thesis (Chapter 2) is the literature review and will present the literature associated with these areas of research.

Chapter 2: Literature Review

2.1. Introduction

Chapter 1 of this thesis has highlighted the background area and motivations behind conducting this research. This chapter will follow on from the introduction and summarises key work undertaken by others in the area and outlines how this has been incorporated into the development of the work undertaken in this PhD. The chapter has been split into several sections, the next section describes the search that was undertaken for the literature review (2.2), the third section discusses home monitoring systems (2.3) and the fourth section specifically describes home electricity monitoring systems (2.4). The fifth section discusses the gaps in research (2.5), which this literature review has found, and how these can be investigated as part of this research.

2.2. Literature review methodology

For this a comprehensive literature search was conducted using multiple databases including Google Scholar, IEEE online xplorer, Scopus, ACM digital library and Medline. Some examples of the search terms used are, ‘smart homes’, ‘health smart homes’, ‘telecare’, ‘telecare technologies’, ‘electricity monitoring’, ‘electricity activity monitoring’. There was no exclusion put on the search criteria but only papers publish in English were accessible for the researcher.

2.3. Home monitoring

2.3.1. Introduction

This section will outline the development of direct monitoring systems, with both passive and active sensors, that can locally, or remotely, monitor and support residents in their own home. The section will then discuss how this has been an evolving area of research, with reference to specific projects and will describe some of the range of sensors developed to support monitoring. The section outlines the placing of more targeted sensors in the care of people with health problems and the placing of sensors in the health care environment (telecare). Linked with this, the review considers methodologies for producing a standardised measure to assess an individual’s capacity to live independently. The section concludes with discussing research that explores ethical issues surrounding the use of monitoring, specifically, the issue that while monitoring of activities can assist independent living, it can become intrusive for the people whom it is intended to support.

Within research into smart homes and health smart homes, as highlighted by the reviews of Chan, Estève, Escriba, & Campo, (2008), Reeder et al., (2013) and Alam, Reaz, & Ali, (2012) there have been many different examples of smart homes (as defined in section 2.3.2) and health smart homes (as defined in section 2.3.3). This section aims to provide an overview of some of the different smart homes, health smart homes and telecare systems that have been developed and utilised by researchers.

2.3.2. Smart homes

The development of fixed and wireless communications, and much faster and more reliable web-based technologies, combined with the decreasing cost of different types of sensing technologies (as described in section 2.3.5) has led to the increased use of managed, sensor-determined support in the home environment (Chan et al., 2008). Homes in which remote sensing has been introduced have been given a collective title of "*smart homes*".

Jiang, Liu, & Yang, (2004, p.659) defined a smart home, as *"a dwelling incorporating a communications network that connects the key electrical appliances and services, and allows them to be remotely controlled, monitored or accessed"*. There has been research using specifically constructed test-bed smart homes with a range of different sensors, as well as systems using existing buildings with the installation of a small number of sensors (typically one or two). Research into smart homes has outlined supporting and developing environments that will improve the resident's comfort, safety and/or wellbeing.

An example of a specially constructed home is The Adaptive House by the University of Colorado (Mozer, 1999,1998). The Adaptive House incorporates a system called ACHE (Adaptive Control of Home Environment), which was developed to meet two objectives; the first was to anticipate the residents' needs and the second was to support effective energy conservation. To achieve these objectives, the system was connected to sensors in each room that monitored the information about the room environment (for example, temperature and light intensity, etc.). In addition to room sensors, the system also received other information about the house, for example, the water heater temperature, energy usage of appliances together with the gas and electricity costs. ACHE continuously monitored the environment and the needs and wants of the residents to learn their lifestyle preferences (for example, the temperature or lighting preferences). From the data gathered from the sensors, the ACHE system then predicted the residents' optimum environment and set the house to their preference in the most energy-efficient way.

Simulation studies were run on the heating control of the house (Mozer, Vidmar, & Dodier, 1997) with the results from these showing that the developed ACHE system performed better than three other non-adaptive controller providing a lower mean daily energy and discomfort cost. The adaptive house and ACHE are one of the earlier examples of work in the area of smart homes and the work in this project was limited to home environment control (for example lighting, heating and temperature control of the home). As discussed in other examples in this literature review, the work of smart homes, health smart homes and related disciplines has expanded to incorporate a much larger number of sensors to monitor a much larger range of activities and is described below.

The MavHome smart home project at the University of Texas at Arlington (Das, Cook, Battacharya, Heierman III, & Lin, 2002) aimed to achieve maximum comfort for those living in the home as well as being energy efficient and minimising running costs. The MavHome used a number of algorithms to predict the residents' movement and interaction throughout the house. The MavHome system was implemented in the MavPad (Youngblood, Cook, & Holder, 2005a) apartment. The MavPad had a number of difference sensors (for example lighting, humidity, temperature, motion sensors etc.) that provided the information to the MavHome system (Youngblood, Cook, & Holder, 2005b). The MavPad was in operation for one year, with three different residents. During this time different observations and experiments were run using the system, the first of these involved testing the sensors installed into the MavPad to see if patterns in the inhabitants activities could be discovered from the sensor data (Youngblood et al., 2005a). The second experiments involved collecting data from an inhabitant on just motion and lighting control with the aim of reducing the inhabitant's interaction with the lighting within the MavPad. This experiment showed a 54.9% reduction in interactions. For the final experiment an individual occupied the MavPad for 9 months, during this time different observations and automations were conducted (Youngblood et al., 2005a). From this experiment the full system managed to automate 39.98% of the inhabitant's life. It was noted by the authors that the inhabitant used for this final experiment led a very erratic lifestyle, which was a challenge for the system to learn (Youngblood et al., 2005a). The authors also highlighted a number of issues, which were observed during this period such as failures of the sensor network and unreliability. These are some of the areas that cause problems with the use of sensor technology to monitor aspects or activities of a person; section 2.3.4 discusses these issues in more detail.

A further test bed was the intelligent dormitory (iDorm) at the University of Essex (Pounds-Cornish & Holmes, 2002), which was a room designed as a student dormitory based on university accommodation. It was seen as an all-inclusive room where the resident would undertake a number of activities, for example, sleeping and working. The iDorm was fitted with a number of sensors (humidity, temperature, light etc) as well as embedded sensors in the furniture (pressure sensors). From the data gathered from the sensors, the iDorm system was tested with the aim of learning the patterns and the needs of the resident and adapting certain features in the room to their needs (for example, setting appropriate temperature and lighting) (Hagras et al., 2004).

The research carried out into smart homes and the use of sensor technology in smart homes is not limited to academic research. Several commercial companies have also carried out research into smart home technologies that can be used as off the shelf products. An example of this is the work by Philips in the development of their HomeLab (De Ruyter & Aarts, 2004) in which they tested some of their new entertainment products. Microsoft has also researched the use of sensors to track and identify people movements around rooms called the EasyLiving Tracker (Krumm et al., 2000). There are also several web-based companies dealing with the supply of technologies used for smart homes (Smarthome, 2015; Smart Home Supplies, n.d.).

The examples above are ones in which the focus has been using sensors to acquire learning, so as to understand the preferences of the resident and to provide them with an optimum environment in the most energy-efficient way. There has also been a focus on developing individual support systems for addressing the health needs of older and frail people. These developments aimed to capture physiological data, through sensors, as well as provide some improvements in comfort for the residence. This approach has been known as health smart homes.

2.3.3. Health smart homes

Noury et al., (2003, p.118) defined a Health Smart Home as a home that allows "*an autonomous life, in their own residence, to people suffering from various pathologies and handicaps, which should normally force them to be hospitalised or placement in specialised structures*". The term Health Smart Home is now applied very widely to almost any home with any sensors that are used to monitor any health condition. This health need may vary from one that requires minimal intervention for the majority of the time, for example, monitoring a person with type 2 diabetes, to a health need that requires almost constant intervention, for example, a patient with severe heart failure. In developing home-based monitoring, the rationale for independent living is not only supporting the wishes of the individual, but also to reduce the potential economic burden of placing the individual in an institutional home or in hospital. Health smart homes can be seen as a test bed for the trial of different types of techniques and sensors to support the monitoring of elderly or disabled people.

The HIS project at the Faculty of Medicine at Grenoble, France (Demongeot et al., 2002) developed an apartment that was fitted with a variety of sensors. These sensors were used to gather both physiological data (for example, heart rate and blood pressure) as well as ambulatory actimetry sensors (that are used to detect physical activity and so are used to detect falls or lack of movement). Additionally, IR (Infra-Red) sensors were placed around the apartment; these were used to detect the movement of the residents around the apartment. The data from the sensors were transmitted to a remote monitoring station and were logged. The logged data were then reviewed and, in case of a danger, for example a fall, an alarm was triggered. To show different functions of this health smart home, different simulation studies have been conducted (LeBellego et al., 2006; Virone, Noury, & Demongeot, 2002).

The monitoring of the residents' physiological data and movement was also used in the PlaceLab project (Intille et al., 2006). In this project, the residents used wearable sensors to monitor movement and vital signs (accelerometers and heart rate sensors). This project also used video cameras and microphones as other sources of monitoring of the residents. The PlaceLab project also outlined the monitoring of the environment (temperature, humidity, etc.) and energy usage in the apartment (gas flow, water flow, etc.). The apartment also included a number of on-off or open and close sensors on appliances so that the opening of the fridge and turning the oven on and off was recorded. A number of pilot studies were carried out at the PlaceLab. These pilot studies were for both observational and collecting sensor data. The captured data allowed for analysis, both in terms of the possibility of predicting the activities of the residents but also to investigate the data collected from sensors.

The Welfare Techno House (Kawarada et al., 1998) also monitored the resident's physiological signals and movement, but only while they were asleep, bathing or using the toilet. Electrocardiogram (ECG) monitors were used to monitor the resident while bathing and sleeping and a body and excrement weights were used to monitor resident while on, and using, the toilet. This project also incorporated environment sensors that were used to control the lighting throughout the house. Several experiments were run in the house to tests the sensors and their effectiveness of collecting information (Tamura et al., 2007; Tamura, Togawa, Ogawa, & Yoda, 1998).

The Gator Tech Smart Home (Helal et al., 2005) differed from the three homes described above in that it did not monitor the residents' physiological signs but was developed to support their cognitive impairment. The home incorporated technology to provide reminders for taking medicine and appointments. It also had many different types of appliances (smart mailbox, smart fridge, smart phone etc.). These smart appliances monitored their usage and sent notification to the occupant, for example, the smart mailbox informed the occupant when the mail had arrived and the smart fridge informed the occupant when food had exceeded its use-by date. The home also had a number of energy and environmental sensors (smart thermostats and smart plugs), security and activity monitoring (home security and an emergency monitoring system).

The CASAS smart apartment at the Washington State University (Helal, Cook, & Schmalz, 2009) was equipped with similar sensing equipment to that of the Gator Tech Smart home (Helal et al., 2005) described above. The CASAS smart apartment incorporated motion, light, temperature, humidity and door contact usage sensors (Helal et al., 2009). This apartment has also been equipped with specific item sensors, for example to detect water and oven usage (Singla, Cook, & Schmitter-Edgecombe, 2008).The aim of the CASAS is to recognise the resident's activities based on sensor data recorded from the test environment (Singla et al., 2008). An experiment was carried out using 22 volunteers, with the aim of using a Markov model algorithm to recognising when the volunteers had performed an pre-defined range of activities, called activities of daily living (ADLs), as described in section 2.3.6. Based on the results from this experiment, the algorithm, a Markov model (as described in section 2.3.7), gave an overall accuracy of 88.63% for recognising the ADLs performed (Singla et al., 2008). Other smart homes, health smart homes, telecare and related disciplines have used data mining and machine learning techniques to recognise when activities or ADLs had been performed from data provided by sensors. This is discussed in more detail in section 2.3.7.

The Aware Home (Kidd et al., 1999), similar to the Gator Tech Smart Home, did not monitor the residents' physiological signs, but consisted of a smart floor that tracked the movements of the occupants and built up patterns of the residents' movements. The Aware Home also provided cognitive assistance in the form of a system for finding the most frequently lost items. For example, keys could be located by attaching a small radio frequency tag to these frequently lost objects.

The Ubiquitous Home (Yamazaki, 2005, 2006, 2007) was a test bed used for the creation and testing of new services for the home by the linking of devices, sensor and appliances through data networks (Yamazaki, 2007). Within the Ubiquitous Home there were a large number of different sensors, which had the aim of monitoring activities. Each room within the Ubiquitous Home had a video camera and a microphone in the ceiling to gather information. The presence of audio and visual recording equipment in this house raises concerns about resident's privacy, which will be discussed later in section 2.3.8. This house incorporated a number of different sensors, which measured the movement of the resident. These included floor pressure sensors that also detected the positions of furniture as well as tracked the movement of the residents. IR sensors were located at the entrance of doors as well as in the corridors and in the kitchen. There were also two radio identification systems (RFID), which detected when people wearing RFID tags entered the rooms. The house also included vibration and accelerometer sensors to monitor movement further. As well as monitoring the residents, the house could also provide certain context-aware services, such as a television (TV) program recommendation service. Several observation and trial evaluation experiments have been run in the Ubiquitous house to test the working of several features of the house (Yamazaki, 2005) as well as to provide feedback on some of the sensors and systems implemented into the house (Yamazaki, 2007).

The U-Health smart home (Agoulmine et al., 2011) aimed to help support elderly or chronically ill people in their own homes. The house incorporated a number of sensors and actuators to monitor and support the resident. The sensors were used to collect environmental data (such as temperature and humidity), as well as physiological information from the resident (such as heart rate). As well as the sensors, the house incorporated a number of actuators, which were used for a number of different tasks such as turning off appliances and lighting. The data collected from these sensors were analysed to highlight any health and safety concerns about the resident, or to perform certain tasks for the resident such as turning off appliances (Kim et al., 2010).

2.3.3.1. Discussion

The ranges of sensors placed in the specifically built homes (as shown in sections 2.3.2 and 2.3.3) have allowed the monitoring of a range of physiological, cognitive and environmental areas. These homes provided a range of effective test beds for developing systems and, in particular, linked sensor activity to patterns or changes in patterns of behaviours. However, these homes also highlight a number of issues such as set-up and running costs of the equipment (Chan, Campo, Estève, & Fourniols, 2009), ethical issues of the use of sensors to monitor (Chan et al., 2008) and the varying needs of the patients (Chan et al., 2008).

A recent review of health smart home technologies by Reeder et al., (2013) again highlighted that the development of sensor-based systems has the potential to support older adults remaining within the community. This is highlighted by the wide range of different technologies developed and utilised by researchers, as described in sections 2.3.2 and 2.3.3. However the review by Chan et al., (2008) showed that the limited size and length of these studies have made it hard to evaluate the technologies in real life, long-term situations (Chan et al., 2009). This is highlighted in the research described in sections 2.3.2 and 2.3.3, with these showing a wide range of technologies used to monitor a wide range of variables, though for only short periods of time. The review of smart home technologies, by Demiris & Hensel, (2008), identified that this field of research was at a relatively early stage of development and acknowledged that there was a lack of evidence to support one approach or another. This is a view supported by the research described in sections 2.3.2 and 2.3.3. The studies highlighted in sections 2.3.2 and 2.3.3 have shown the feasibility of using different sensor technologies, in different scenarios or experiments though with a limited number of people. These studies also provide no indication if the use of a certain type of monitoring technology is better than another type or how the developed technologies could be adopted in wide-scale real-life situations.

As outlined previously, and highlighted in reviews of Alam et al., (2012), Chan et al., (2009, 2008) and Demiris & Hensel, (2008) another important consideration with the use of home based remote sensors is privacy and personal independence of the resident. This issue is discussed in more detail in section 2.3.8.

2.3.4. Telecare

The work of Tang & Venables, (2000, p.8) referred to telecare as "*remotely delivered care and support; this might include rapid response to emergencies in home, treatment and medical advice, and continual monitoring*". Telecare, therefore, aims to deliver support to people in their own homes to be used alongside existing care packages (Stowe & Harding, 2010). The work of Stowe & Harding, (2010) also highlighted potential benefits to the resident, as well as for the caregivers and society more generally, from the installation of different telecare systems. Some of these potential benefits include a reduction in unnecessary hospital admissions, a reduction in falls for the person and an increase in personal independence.

Within the area of telecare there are also two other terms telehealth and telemedicine are sometimes used interchangeably (Barlow, Singh, Bayer, & Curry, 2007). Telehealth is described as the management of long-term health conditions through monitoring (Stowe & Harding, 2010), whereas telemedicine is the use of different communication technologies to provide healthcare at a distance (Koch, 2006; Stowe & Harding, 2010). For this literature review, the research is mainly focused on telecare and the different technologies employed to monitor and support people within their own homes.

Within telecare, there are three generations of systems that can be identified (Brownsell et al., 2008; Stowe & Harding, 2010). Table 2.1 gives an overview of each of these generations and examples of the types of systems or sensors, which are classed within each generation.

Telecare Generation	Explanation	Examples
1st Generation	Systems that require the user to activate the system if they require assistance	Emergency assistance alarms worn round the resident's neck
2nd Generation	Systems that can automatically detect emergencies without the user input	Fall sensors
3rd Generations	Systems that are used to monitor resident's activities or behaviour so as to highlight changes in their behaviour that can be explained by changes in the residents health. (Brownsell, Bradley, Blackburn, Cardinaux, & Hawley, 2011)	Lifestyle monitoring

Table 2.1: Description of the different telecare generations

The basis of lifestyle or behavioural monitoring is to use sensors placed in and around the home to monitor a set of everyday activities undertaken by the resident and then to identify changes from the data collected that could then be used to suggest a deterioration in the health of the resident (Hanson et al., 2007). For this type of monitoring, a set of activities needs to be chosen as a measure that can be recorded in respect of time of duration and frequency of repetition, with regular recording of these highlighting changes over a period in how each or all of the activities are being performed. For choosing the activities used in lifestyle monitoring, there are two approaches. The first is to use a pre-defined range of activities, called activities of daily living (ADLs), which are described further in section 2.3.6. The second is to use more general activities, either in a collective group of many activities, or as individual activities that are deemed necessary for living and could be used to show a change in the user's lifestyle that could indicate a health change. Below, two examples of lifestyle or behavioural monitoring systems are discussed.

Alwan et al., (2006) used a variety of different sensors to monitor ADLs for 22 residents (15 of whom did not have dementia). The study work was undertaken over 3 months and was confined to residents of the Homestead at Maplewood in North America and involved 7 males and 15 females with a mean age of 83.8 years (only one resident was under 65). Residents were excluded from the research if they declined to be included or if they required extensive outside assistance for daily living or were unable to get out of bed unaided. Monitoring was achieved by placing a number of different types of sensors around the resident's home; these sensors included passive IR sensors in every room as well as in the shower, a stove top temperature sensor and a bed occupancy sensor.

As well as monitoring of ADLs, the system developed by Alwan et al., (2006) was also set up to give extra key alerts, for example, that the resident had left the stove on if they had gone back to bed or left the kitchen for more than an hour. The data from this system were analysed and the caregivers were given access to summaries for each of their patients with an accompanying score based on their ADL performance. The system also had the potential of sending out alerts directly to the caregiver if it was deemed necessary. In reporting the results, the authors identified cases in which the data may have assisted early diagnosis, including rectal bleeding, where the monitoring detected increased visits to the bathroom (from four times daily to 15 times daily), a patient who showed increased restlessness which was resolved with increased pain control, and a patient with congestive heart failure who was sleeping poorly, which the monitoring detected, and which was resolved by raising the head of the bed.

In this study, Alwan et al., (2006) showed that there was acceptance of the use of monitoring both by the residents and the formal and informal caregivers. The system did produce false alerts, however with further investigation, many of these were related to the technology malfunctioning, or being tampered with, and it was suggested that this could be significantly reduced with better setting up of the system. The system did show the potential of using sensors to monitor a predefined list of activities (ADLs) to highlight changes or issues in residents' health or social care status and the benefits as a way of coordinating the care given to each patient more effectively, depending on their needs.

In comparison, Glascock & Kutzik, (2006, 2007) developed a behavioural monitoring system that focused on monitoring general activities rather than the predefined ADLs. Glascock & Kutzik, (2007) developed a system that monitored five general activities, i.e., wake-up time, meal preparation, medication adherence, overnight toileting and general activities levels during the day. This system worked by placing a number of passive IR sensors in strategic places throughout the resident's home and near key appliances. For example, an IR sensor was placed in the medication box to show that they were using this (i.e., taking their medicine). A series of studies was done across different sites in Britain and America with a total of about 350 residents using the system. The patients from different studies had a range of different clinical needs, frailty and social needs and were not as pre-selected as those in the study by Alwan et al., (2006).

The information collected from Glascock & Kutzik's (2006, 2007) system was used to identify a summary of activities that the resident had performed that day. The activities were shown next to a traffic light system with green indicating that the activity was performed with no significant deviation in frequency or timing from the previous data and red and yellow indicating when there had been a significant deviation in frequency or timing above the thresholds set for the two markers. This information was then provided to the caregiver for them to assess what the next course of action should be, for example, a home visit or telephone call. The system did have a significant number of false alerts, some of which resulted in the emergency service forcing entry into the accommodation whilst the resident was out shopping. However, the system did also detect, at an early stage, some potentially serious events. For example, the system detected a dramatic increase in night-time use of the bathroom in a 49-year-old human immunodeficiency virus (HIV) patient. On investigation the patient was found to have hives and was rapidly treated and the situation quickly resolved. Another case showed a resident who, over a number of days, remained in their home (against previous patterns of activity) and was visiting the bathroom over 25 times daily. The patient was found to have a urinary tract infection that was quickly managed with antibiotics.

As outlined, both of these systems and other work in the areas of lifestyle monitoring have highlighted the issue (Brownsell et al., 2008) of these systems producing a large number of false positive alerts, although in both the work of Alwan et al., (2006) and Glascock & Kutzik, (2006, 2007) it was highlighted that the number of false alerts was in part due to technology issues. For these systems to be effective in their recognition of changes in behaviour patterns, it is essential to minimise false alerts so as to not overload the caregivers with incorrect information, or worse, so that they no longer trust the system when it was telling them there was a genuine problem with one of their patients.

ADLs were developed as an assessment scale to assess whether the patient was capable of living independently or that they required assistance (as discussed in more detail in section 2.3.6). These assessment tools and other similar ones were developed to be used by health care professional and caregivers to carry out the assessments on the patients in person. Research has now been extended, as shown above, to use sensors to measure a patient's ADLs. This carries its owns problems of intrusiveness as well as cost (Glascock & Kutzik, 2006, 2007) and the accuracy of such systems in detecting potential problems.

2.3.5. Sensors

As discussed in the sections above, the use of many different types of sensors have made the collection of data possible within smart homes, health smart homes and telecare systems. Sensors can be used to monitor a large number of different aspects about a person's life, their actions and their environment. Sensors are also very versatile and can take many different shapes and forms. The following sections discuss how different sensors are used to monitor many different aspects of the resident and the environment.

2.3.5.1. Wearable sensors

Sensors come in many different forms and can either be fixed to the house structure or objects in the house or worn by the resident. Worn sensors can either be fitted into clothing, for example shirts with sensors sewn in (Paradiso, 2003; Park & Jayaraman, 2003) or sensors that can be worn around the wrist (Anliker et al., 2004). These wearable sensors are used to monitor the patient's physiological data (such as blood pressure and heart rate) with very little or no operational input needed from the patient. The health smart homes of Agoulmine et al., (2011), Demongeot et al., (2002), Intille et al., (2006) and LeBellego et al., (2006) also incorporated wearable sensors with other fixed sensors within their health smart home environments.

2.3.5.2. Sensors for movement

The monitoring of the movement of the residents around their homes has been widely monitored within smart homes, health smart homes and telecare. There are many different types of sensor used to monitor movement, from floor pressure (Helal et al., 2005; Kidd et al., 1999; Yamazaki, 2005, 2006, 2007) to IR sensors (Alwan et al., 2006; Glascock & Kutzik, 2006, 2007).

2.3.5.3. Sensors placed on objects

Another group of monitoring sensors is ones that have been placed into objects or appliances to monitor their use specifically, or the resident while they are there. Examples of sensors used to monitor the occupants are sensors placed into the bed to monitor heart rate (Kawarada et al., 1998) or pressure sensors on chairs and beds to monitor occupancy (Brownsell et al., 2008; Hanson et al., 2007).

Sensors can also be attached to appliances or objects, to monitor usage. Examples of these are sensors placed into a cooker to monitor whether the cooker has been left on (Adlam et al., 2004), sensors placed into a box to monitor if medication has been taken (Glascock & Kutzik, 2006, 2007) or sensors attached to electrical appliances to monitor energy usage, for example, gas, electricity and water flow, within a building (Intille et al., 2006) or the monitoring of the usage of specific electrical appliances (Helal et al., 2005).

2.3.5.4. Sensors to monitor the environment

As well as sensors to monitor information about the resident and their activities, sensors have been used to monitor information about the environment. For example, in Mozer, (1998, 1999) and Intille et al., (2006) sensors were used to monitor the environmental condition inside the house, e.g., the temperature, humidity etc.. Sensors can also be used to monitor certain aspects about the environment that could cause a hazard for the resident, for example, flood detectors or intruder alarms (Brownsell et al., 2008).

2.3.5.5. Sensor summary

This section has highlighted the many different forms that sensors can take and the many different types of activities sensors can be used for monitoring. There are, however, a number of privacy and security issues surrounding the use of sensors to monitor activities; this is discussed in more detail in sections 2.3.8 and 2.3.9. Section 2.3.7 will discuss the many different machine learning and data mining methods used to analysis the vast amount of data, which are generated from these different sensors and how it may be interpreted.

2.3.6. Activities of daily living

Much of the health based research has been focused towards the range and levels of support to retain independent living and Katz, Downs, Cash, & Grotz, (1970) introduced a tool - The Index of Activities of Daily Living to provide standard measures which have been used widely in the intervening years. The tool measures the performance of the elderly or disabled person to be able to live independently by summarising their ability to complete certain activities, for example, washing themselves, cooking, eating and drinking, based on a three point scale, indicating if they required no assistance, some assistance or complete assistance to perform each activity. Katz et al., (1970) developed their index of ADLs by observing activities being performed by a group of patients each of whom had had a fractured hip. From this work, other scales have been developed that use the ability to perform other functions such as answering the telephone and doing the laundry (Lawton & Brody, 1969). Other ADL scales have been developed to test those with specific diseases or illness, for example, dementia (Bucks, Ashworth, Wilcock, & Siegfried, 1996).

Section 2.3.7 gives examples of research that has been undertaken to automatically recognise or classify data given from sensors monitoring ADLs or other activities. Other examples of the recognition of ADLs or activities from electricity monitoring are discussed in section 2.4.

2.3.7. Data mining and machine learning for activity recognition

Research has been undertaken to automatically recognise data generated using sensors either relating to ADLs or other activities. This is used as a way to monitor and highlight changes in the ADLs or activities or the resident(s). To recognise the activities from the data given by the sensors, different data mining and machine learning techniques have been utilised to transform the outputs from the sensors into the activities performed.

An example of this has already been given in the work of Singla et al., (2008) which utilised a Markov model to recognise the performing of ADLs from sensors data collected from a smart home environment. Fleury, Vacher, & Noury, (2010) used a support vector machine (SVM) to classify the data given from multiple sensors into specific ADL activities automatically. An experiment was conducted using 13 participants at the health smart home at the Faculty of Medicine at Grenoble, France. Each of the participants was asked to perform a list of ADLs and the data collected from this experiment were used to construct a SVM. The SVM was then used to classify the data into ADL activities performed. From this classification the SVM gave a mean overall accuracy for all of the activities of 86%. It was noted by the researchers that all the participants were young and healthy and any further studies should incorporate elderly people living independently at home and be for a longer time period, as the time period for the collection of data for this experiment was short, with a mean time of 51 minutes. Other examples include the work of Tapia, Intille, & Larson, (2004) whose research used a Naïve Bayes classifier to recognise a list of ADLs, using a large number of sensors placed on objects around the home. In addition, Kröse, Van Kasteren, Gibson, & Van den Dool, (2008) used a Hidden Markov Model, to recognise a list of ADLs from a number of different sensors placed into a participant's home.

As well as the recognition of ADLs, machine learning and data mining techniques have also been used to recognise activities, which were predefined by the researcher. Yang, Wang, & Chen's, (2008) research used an artificial neural network to recognise a list of eight activities from data collected using an accelerometer worn by the seven participants who participated in the study. Some of these activities included walking, running, standing and vacuuming. The best overall accuracy from this experiment was 95%, although it was noted by the researchers that they aimed to use this system to recognise more complex activities.

This section gives some examples of the different data mining and machine learning techniques that have been used to analyse sensor data, from both fixed and wearable sensors, to recognise activities. Section 2.4.5 will then discuss in more detail data mining, machine learning and pattern matching, as well as the methods used in each of these. Chapter 3 of this thesis discusses in more detail the steps undertaken in the development of model or algorithms used to recognise activities for this research in this thesis.

2.3.8. Privacy issues with the use of sensors

As can be seen from the examples of work given in the previous section, all of these projects incorporate a wide variety of sensors that are used together to monitor many different aspects of a resident's pattern of living. The use of all these sensors has raised concerns about privacy of the data collected as well as the intrusiveness of the sensors particularly in respect of long-term collection of data.

Some of the terms that are used to describe the tasks performed by different sensors can have negative connotation. Perry et al., (2009) give examples of terms such as surveillance and tagging, which can be associated with criminal behaviour. Landau, Werner, Auslander, Shoval, & Heinik, (2009) highlighted technologies, such as global positioning system (GPS) tracking, which have been used to monitor dementia sufferers in case of wandering. This is the same technology as used by civic authorities to monitor criminals (Stowe & Harding, 2010). Although the intentions of the surveillance are different, for example surveillance of people with their consent compared to surveillance without people's knowledge or consent (i.e., Government or police surveillance), the principles of them both are the same (Perry et al., 2009). Therefore, care needs to be taken with the monitoring of individuals to ensure consent is obtained or is seen to be in their best interests.

The perceived intrusiveness of a system, which can also highlight the reduction of privacy, can be a key factor in the acceptability of a monitoring system (Bowes, Dawson, & Bell, 2012). Demiris, Hensel, Skubic, & Rantz (2008) conducted focus groups with elderly people to understand their views on the use of sensors in their homes. The participants in this study gave their views on certain types of sensors. Fall sensors were one of the sensors that were uniformly supported by all the participants, motion detectors were also widely supported, but for the use in the detection of intruders, not for the detection of activities of the inhabitant themselves. Most participants did not support other sensors (for example, video cameras) as they raised the most concerns about privacy. Sixsmith et al., (2007) also highlighted that the use of technologies, such as video and microphones, can be seen as very intrusive. Stowe & Harding, (2010) also highlighted the issues with the installation of different technologies and how their visual presence, or having to use the technology specifically for this system to work, had an effect on its uptake. There are also concerns that the presence of such equipment can leave individuals embarrassed about their vulnerabilities and remind them of hospital or institutional environments (Stowe & Harding, 2010). A way to overcome these issues would be to develop an unobtrusive system that would collect data without the user being aware of data being collected (Bowes et al., 2012).

The instillation of systems that collect information about the resident without their conscious knowledge has a way of minimising the intrusiveness issues around telecare. However, this could then be seen as “covert” surveillance (Stowe & Harding, 2010) even with the consent of the occupant. The work of Demiris et al., (2008) also noted that participants of the focus groups wanted the technology to detect emergencies and not to monitor or detect trends in their behaviour.

With the installation of monitoring technologies in individuals' homes there is the issue that the ability to ask for help has been taken out of the hands of the individual (Bowes et al., 2012). There is also the fear that if something were to emerge from these systems the caregivers and family might then think that the individual could no longer look after themselves and may force them into accepting institutional care (Percival & Hanson, 2006).

Another perceived problem with the use of monitoring technologies was that this could lead to a reduction in day-to-day contact with their care providers (Chan et al., 2008; Draper & Sorell, 2013; Milligan, Roberts, & Mort, 2011; Stowe & Harding, 2010). The work of Bowes et al., (2012) gave an example based on a lifestyle monitoring system, that if the caregivers or family were reassured that their family member was well, that could affect their contact with them. The perceived reduction in contact with care professionals is widely mentioned in the literature, and the work of Milligan et al., (2011) highlighted that any use of a monitoring system should be used as an addition to the giving of care, not a replacement for it.

Gaining informed consent from the residents involved in remote monitoring is a fundamental ethical requirement (Bowes et al., 2012). The participants must be fully aware of what data are being collected and to what end the data will be used. The issue of confidentially, similar to privacy, is also an issue in the use of telecare systems, with some perceptions from individuals that everyone can “see” into their homes (Magnusson & Hanson, 2003) or not wanting particular people having access to their data (Milligan et al., 2011). This highlights the issues of data security (as discussed in section 2.3.9), as well as the importance of keeping the data confidential and only being disclosed to those whom the user has agreed to see it. These are also a very important priority of any monitoring system being installed.

With some monitoring systems, there is an issue of getting informed consent from the occupant when the researchers are not sure what their system will reveal (Bowes et al., 2012). Bowes et al.'s, (2012) example was of a lifestyle monitoring system which was used to monitor changes in activities, although without being specific about what were those changes. This then raises an ethical issue about collecting data for an end that the individual is not informed about (Bowes et al., 2012).

The second issue with gathering informed consent is the case of those with cognitive problems. In these cases, the individual might not be fully aware of what data are being collected and why (Bowes et al., 2012). However, for caregivers and families, the installation of these systems can provide reassurance even if the systems are intrusive, for example, the use of GPS trackers in case of wandering by people with dementia (Landau et al., 2009): the use of these types of monitoring could be seen as “covertly spying” on a relative (Perry et al., 2009).

Within the design and implementation of a monitoring system, there needs to be the ability to turn parts, or all, of it off (Stowe & Harding, 2010). This is an important ethical point, as it could be argued that the user must have the right to withdraw from the monitoring, should they so wish.

The implementation of these types of monitoring systems can have an effect on the privacy of the individual, although the benefits, as well as the way they are designed, can minimise this reduction in privacy. As already discussed in this section, there are several design considerations that must be followed before developing and implementing a monitoring system. The first of these is data security; the data collected by sensors must be secure. Furthermore, the data collected must be confidential and only disclosed to those whom the user decides. The perception of intrusiveness is also an important consideration when designing a system, as the more non-intrusive the system is and the greater the ability for it to collect information without the direct input of the user, the better its uptake. The aims of a monitoring system need to be clearly defined and explained to the users. Finally, the system must be able to be turned off by the user should they so wish.

2.3.9. Security issues of sensors

As well as privacy and ethical issues of the use of sensors to monitor people, there are also a number of security issues with the use of sensors. The works of Ameen, Liu, & Kwak, (2012), Chan & Perrig, (2003), Kargl, Lawrence, Fischer, & Lim, (2008) and Meingast, Roosta, & Sastry, (2006) have highlighted the issue of security of sensor data used for health purposes or for general purposes. With the large amount of data being shared electronically, there becomes a greater risk of attacks from hackers and other malicious attackers, than if the data was in paper form (Meingast et al., 2006).

The work of Ameen et al., (2012), Chan & Perrig, (2003), Kargl et al., (2008) and Meingast et al., (2006) have also highlighted other ways in which sensor networks can be attacked and used for different means. One example of this is eavesdropping, where the data from the sensors is stolen and used for criminal or malicious purposes. The use of eavesdropping can also put the safety of the occupants at risk, as certain information about the (vulnerable) occupant can be accessed through this process, for example their location (Ameen et al., 2012).

The problem with using sensors is not that they are used to collect information. As Chan & Perrig, (2003) argued, almost the same information can be collected from direct surveillance of the house as with sensors. It is the scale that is the real problem. As the installation of numerous sensor networks means that large amounts of data from different places are being generated, it can be hacked and monitored from one location. As discussed in section 2.3.4, this is however a clear benefit of remote monitoring with sensors, that multiple people can be monitored from one location, provided that the data are collected and transmitted securely.

Meingast et al., (2006) and Ameen et al., (2012) discussed a number of different methods that can be implemented to ensure the security of the data collected from sensors. An example of this is data encryption, so that the data that are collected from these sensors are encrypted therefore making it harder for hackers and others to read the data.

2.3.10. Summary

There has been extensive development in the use of home-based monitoring and sensing equipment as shown by the range and breadth of sensors from a number of test-bed homes. This development has been driven by the reduced acquisition cost of sensors, an increased reliability of sensors not to give false triggers and an improvement in the transmission of sensor data from remote monitoring systems. The placement of sensors has been explored in both improving environmental comfort of the resident and also in supporting the residents to maintain independent living. However, as shown by the work of Demiris et al., (2004, 2008) people may be reluctant to have a range of sensors monitoring their daily actions and may actively seek to undermine such support systems by for example not wearing the equipment (Demiris et al., 2004).

The development of sensors has allowed many to become much less intrusive; however, they still require installation and set-up investment. The use of ADLs has been effectively linked with the use of sensors so that if the resident accepts the intrusive nature of sensors this can provide effective daily monitoring.

In summary, this section has highlighted there has been a significant range of research in monitoring of residents using a wide array of different sensors. The review shows that such sensors have limitations, with many being perceived as intrusive, some requiring the residents to wear a device, and some requiring the resident to interact with the device. This can lead to residents not wanting to install devices actively or to use such sensors for long-term monitoring. This section has also considered the development of a list of daily task activities to a standardised scale that can be used to monitor a resident's ability to manage to live independently. By monitoring changes in how a resident is able to perform these tasks, they can be used to show changes in the resident's ability to live independently.

The following section will go on to discuss electricity usage monitoring devices and how they can be used to show other information about the occupant and not just the household's electricity consumption.

2.4. Electricity analysis

This section will discuss the different devices that can be installed into homes to monitor energy usage, which provides feedback to the residents. The section will then discuss how data that have been gathered by energy monitoring devices can be used to show activities that the resident has undertaken, based on the appliances that have been used. Finally, the section will discuss how identification of appliance usage may be used to provide the resident with information both about how to use appliances cost effectively but may also be used as a system of remote monitoring.

As well as electricity monitoring, this section will also briefly discuss other types of energy monitoring, which were not included in this research but could be incorporated for further work.

2.4.1. Electricity monitoring

The cost associated with specific energy usage within a home is almost invisible to the user (Petersen, Steele, & Wilkerson, 2009) meaning there is little association between using an appliance and the cost involved. However recently there has been growing interest in the need to reduce energy consumption both due to increased costs as well as to encourage sustainability (Chetty et al., 2008). This has been combined with an increased environmental focus to reduce the carbon footprint and to reduce energy waste (Climate Change Act 2008). One way of achieving these goals is to provide timely feedback on energy consumption, its cost and its environmental cost to the user (Fischer, 2008).

One way of providing feedback to the user about their energy usage, cost and the environmental costs is the installation of a simple monitoring device. These monitoring devices can provide real-time feedback of the current energy usage, which allow the users to see their current energy usage and cost. For energy monitoring there are currently three types of electricity monitoring devices that can be used to provide feedback on electricity usage. These three devices are:

1. Plug socket monitors
2. Whole house monitors
3. Whole house monitors with plug sockets

As well as electricity consumption monitors, there are also gas consumption monitors which can be used in combination with the electricity monitors to provide feedback on the total usage and costs of energy for a home. Section 2.4.1.4, will discuss gas consumption monitors in more detail.

2.4.1.1. Plug socket monitors

Plug socket monitors are devices that can be used to monitor the electricity usage of a single socket. These devices can be used to show the electricity consumption, as well as cost of single or multiple appliances depending on what is plugged into the socket.

Examples of some of these devices are the Kill-A-Watt (P3 International, 2015), Efergy Energy monitoring socket 2.0 (Efergy Technologies Limited, 2014a) and the Belkin Conserve Insight (Belkin International Inc., 2012). These devices provide real time feedback on the energy usage through screens attached to the monitors. As well as the energy usage (in Watts), these devices can also provide feedback on information such as cost, current, voltage as well as the carbon footprint, based on the power usage, of the appliance(s) plugged in.

These devices are useful as they give direct feedback on energy usage and cost of specific appliance around the home. They are also portable so can be plugged into any sock and can provide feedback on any plug-based appliance in the home.

The drawback of these devices is that they can only be used for appliances connected through plug sockets and cannot be used for hardwired appliances, i.e., that are directly wired into the house circuit, e.g., an electric oven. The appliances that are directly wired are generally those that draw the most current and therefore likely to be the most expensive to run (for example, immersion heaters, cooker, shower). These monitors can only be placed on single sockets so they do not provide a full overview of all the appliances within the home. To achieve a breakdown of every appliance, a monitor can be placed on every appliance or socket used within the home, although this can be very expensive (Kulkarni et al., 2011).

The WattBot project (Petersen et al., 2009) used a different approach to overcome this issue of hardwired appliances. In this project, a clamp sensor was placed around each circuit breaker in the home's fuse box to measure the amount of electricity passing through each circuit of the house. This allowed the resident to see the electricity usage of the hard wire appliances in the home, such as the immersion heater, as well as a breakdown for different rooms or appliances, depending on how the home was wired. The recorded electricity data were stored and then downloaded to a mobile device, giving a breakdown of electricity used by each circuit, as well as the total electricity usage for the day.

2.4.1.2. Whole house monitors

Whole house monitors are devices that measure the current electricity usage of the whole house. These devices work by either placing a sensor around the live wire that runs between the electricity meter and the fuse box of the house or by placing a sensor over the light-emitting diode (LED) light on the electricity meter (NB., This is not compatible with all electricity meters due to some meters not having an LED light). These sensors require no direct wiring into the mains electricity of the house. The sensor transmits the real time data to a display unit. The rate at which the sensors capture the current energy usage can vary between models, though for the examples given below the data is captured every 6-12 seconds.

Examples of some of these whole house energy monitors are the Efergy e2 classic (Efergy Technologies Limited, 2014b) and The OWL (The OWL, 2015). These devices come with a display unit, which gives the current energy usage, in Watts, of the house as well as the cost. Some monitors can also be used to monitor environmental features of the house, in the vicinity of where the display unit is located, such as humidity and temperature. As well as being able to see current energy usage, these devices can also be used to save the electricity consumption data and give the option to view historical data via a personal computer (PC) or the Internet.

These devices are generally very easy and quick to install, although there can be problems depending how and where the meter has been installed and the space between and around the live wire. These forms of monitoring also require the user to have access to their electricity meter and for the electricity meter sensor to be within a certain distance of the display unit. These requirements make this type of monitoring more complicated in large houses and/or flats, where the user might not have access to their meter and/or the sensor is placed outside the operating range of the display unit. This could limit their usefulness for monitoring electricity usage for older people living in shared buildings.

2.4.1.3. Whole house monitors with sockets

The third group of monitors are a hybrid group, which combine the monitoring of the electricity usage for the whole house with specific monitoring of a number of single sockets and their attached appliances.

These monitors incorporate the features of the two previous forms of monitoring and so provide the user with the whole house energy consumption as well as the energy consumption of a number of individual appliances. Examples of some of these monitors are the Current cost EnviR (Current Cost, 2015) or the Green Energy Options Ensemble (Green Energy Options Ltd, n.d.). As with the whole house monitors (section 2.4.1.2), these monitors provide real-time information about the current energy usage and cost of the whole house as well as individual appliances. These monitors also come with the option to save the recorded data as well as to view it on a PC or over the Internet. As with the whole house monitors (section 2.4.1.2) the rate at which the sensor captures the information varies by monitor.

These monitoring devices can provide a comprehensive breakdown of energy usage within a home based on appliances as well as the whole house. These devices can also be used to highlight when certain appliances within a home have been used and can also provide other information such as the time the appliance was turned on and the duration of use.

As with the whole house monitors (section 2.4.1.2), these devices are easy to install. However, and similarly, they require the user to have access to the meter and for this to be within a certain distance of the display unit. As with the plug socket monitors (section 2.4.1.1), the plug socket monitors for these monitors can also only be placed on appliances which have plugs to plug into. Although these monitors also record the electricity consumption of the whole house it is possible to break down the energy consumption data to see the energy usage of appliances that are hard wired, for example, the shower.

2.4.1.4. Gas monitoring

As well as monitoring electricity usage, the gas usage of a home can also be monitored to provide feedback to the user, although the ability to install a gas monitor can be more problematic than that of an electricity monitor as they are currently only compatible with certain models of gas meter. The monitoring of gas usage can be useful in determining the overall energy usage of house, although as much fewer appliances within a house use gas, these data provide less granularity compared to electricity usage data.

For the research described in this thesis, the monitoring of a houses gas consumption was not considered due to compatibility problems with meters as well as electricity consumption data providing sufficient depth in the data.

2.4.1.5. Summary

The sections above provide an overview of the different ways in which monitoring devices can be used to provide feedback to the user on their energy consumption, either appliance-based or for the whole house. The use of electricity consumption data to observe appliance usage as an option to monitor residents' activities and their patterns is a growing area of research and will be discussed further in section 2.4.4 of this thesis.

2.4.2. Appliance and activity recognition

The use of electricity consumption data has been a growing area of research to monitor the user's appliance usage for energy saving, appliance recognition or health-monitoring purposes. Although these methods all have different outcomes, the process involved is similar. Each of these methods involves disaggregating electricity data into appliance usage. The results are then used to recognise when an appliance has been used in the home. The data can then provide feedback on ways the user may reduce their energy usage or to monitor appliance usage remotely in order to provide remote feedback about the user's activities and, with analysis, to highlight changes in their activities. Section 2.4.3 will discuss the different research that has been undertaken in the areas of appliance recognition with section 2.4.4 then discussing how this has been extended to the areas of activity recognition.

2.4.3. Appliance recognition

This section will highlight the different approaches that have been used to recognise the use of specific appliances from electricity usage data. The first of these approaches is an area of research called non-intrusive appliance load monitoring.

2.4.3.1. Non-intrusive appliance load monitoring

Non-intrusive appliance load monitoring has been widely used by researchers, as discussed in the review by Zeifman & Roth, (2011), to recognise appliance usage from electricity data. Further examples of how non-intrusive appliance load monitoring has been used to recognise appliances for further analysis include, for example, activity recognition (Belley, Gaboury, Bouchard, & Bouzouane, 2013, 2014), appliance usage monitoring (Rahimi, Chan, & Goubran, 2011, 2012) and heating prediction and control (Spiegel & Albayrak, 2014).

The work into the area of non-intrusive appliance load monitoring was started by the work of Hart, (1992). Hart, (1992) suggested that household appliances could be classed into four separate categories:

1. *Continuous Appliances*: these are appliances that are continuously on and have no or very little change in their power drawn. Examples of

these are fire alarms or clocks. By their nature, many of these appliances draw little power and may be of limited value in monitoring user activity.

2. *ON/OFF Appliances*: This group covers most household appliances that exist in two states, for example, toasters, kettles, light bulbs etc.
3. *Finite State Machine*: These are appliances that have different modes of operation. Examples of these are washing machines, tumble dryers, freezers etc. Whilst these appliances can draw different levels of power, and these levels vary depending on the user's choice of settings, there are only a fixed number of options available. Although this number may be fixed, the number of possibilities when factoring different settings of temperature, spin and cycle, in an example of the washing machine could make the range of these possibilities almost infinite.
4. *Continuously Variable*: This group covers appliances that have varying ranges of operational power levels. Examples of these devices are dimmer lights or power drills. This group can, within a certain range, draw many different levels of power depending on the user's preferences.

Hart, (1992) developed a non-intrusive appliance load monitoring algorithm that was then used to determine the energy consumption of different appliances. This algorithm was based on collecting data from the household electricity monitor (current and voltage at a sampling rate) and then passing it through an edge detector. The detector identified the times and size of changes in the power level. Similar power changes were then clustered together into separate appliances. From this, positive and negative clusters of appliances that exhibit similar size changes were clustered together and then classed as ON/OFF events. Finally, the characteristics shown by the appliances (change in power levels, size of power drawn) were matched against known characteristics of appliances.

The use of this algorithm on experiments in people's homes found that the non-intrusive load algorithm recorded the energy consumption to within +/-10% of the actual appliance energy usage. The experiment also highlighted a number of issues. Although this algorithm was effective in the detection of ON/OFF appliances events, it was less effective at determining which appliances were used within the finite state appliances group. The non-intrusive appliance load monitoring algorithm was also less effective in monitoring:

- Appliances with small power draws;
- Appliances that are continuously on or continuously changing (Appliance type 1 and 4 above);
- At distinguishing between two appliances with the same power draw.

The work in the area of non-intrusive appliance load monitoring has evolved since the work of Hart, (1992) supported by the technological changes in the monitoring devices used for non-intrusive appliance load monitoring. The review by Zeifman & Roth, (2011), highlighted that the areas of non-intrusive appliance load monitoring could be split into two types, low frequency and high frequency monitoring. The monitoring used in the work of Hart, (1992) could be described as low frequency with a monitoring rate of 5 seconds. The review of Zeifman & Roth, (2011) provided some examples of low frequency monitoring devices (with a typical monitoring frequency for a low frequency device being 1Hz). High frequency monitors can be used to monitor at a rate between 10-100MHz (Zoha, Gluhak, Imran, & Rajasegarar, 2012). As described by the work of Zoha et al., (2012), high frequency devices are usually custom made and therefore can be expensive and intrusive to install.

2.4.3.2. Appliance recognition- other methods

Another approach to appliance recognition is ViridiScope (Kim, Schmid, Charbiwala, & Srivastava, 2009). This system involved a combination of different types of sensors. It used an electricity monitor (as discussed in section 2.4.1.2) to record the whole house electricity usage. It also used magnetic, acoustic and light sensors placed on, or near, the appliances to be monitored (Kim, Schmid, Srivastava, & Wang, 2009). Several different types of experiments were conducted to test and validate the Viridiscope approach. The experiments ranged from using just magnetic sensors placed on the cables of each appliance, to the use of light sensors to monitor, for example, a table lamp. The different experiments showed that the ViridiScope system was able to track the power consumption of each appliance tested to within 10% of its actual power consumption. The ViridiScope system could be used to monitor appliances with multiple states as well as appliances with variable power consumption. This system can be seen as a more intrusive form of monitoring, as it required multiple sensors to be placed around the home. Kim, Schmid, Srivastava, et al., (2009) also highlighted that the fusion of all the various sensors as well as their complex installation and maintained needs to be considered when undertaking research in this area. These are highly visible to the user and will provide a constant reminder that they are being monitored.

Patel, Robertson, Kientz, Reynolds, & Abowd (2007) developed another approach to appliance recognition, by using a device that monitored the electrical noise in power lines. The custom monitor can be plugged into any plug socket in the house and uses the notion that each appliance generates a unique noise signature on the power line when turned on or off. An SVM was trained to learn these unique signatures. Once the SVM had been trained it could be used to recognise when appliances were turned on or off. Several experiments were run in a number of houses, this produced a varying overall accuracy of the system of 80-92% for being able to recognise when an appliance was turned on or off.

Unlike other methods, this method did not record the electricity consumption of the appliances. In addition, to recognise the signatures of the appliance it needed to be plugged into a socket in a similar location to that used when the SVM was trained. For example, the signature when connected to a wall socket would be different to the signature when plugged into an extension lead. This work was recently extended by the work of Gupta, Reynolds, & Patel, (2010) with the aim of addressing some of the problems with the original system.

Farinaccio & Zmeureanu, (1999) aimed to identify specific appliances from a whole house electricity data stream. This approach was to monitor the specific appliances that were of interest (for their study this was the water heater and the refrigerator) to capture their electricity signatures. These signatures were then translated into a set of pattern recognition rules relating to each appliance used in the experiment. The development of pattern recognition rules had a number of stages. The first was to detect when the appliance was turned on and off. The second stage consisted of a set of rules that calculated the appliance’s demand profile (i.e., the length of time it had been active). The final stage was to estimate energy usage from the data collected. The two appliances each had a unique set of rules. An experiment was then carried out using a whole house electricity monitor clamped over the main electricity feed into the house. The whole electricity data was then passed through each of the appliance pattern recognition rules. The results for this experiment were given as the differences between the actual daily cost of running the two appliances and the estimated daily cost estimated from the pattern recognition rules. There was a 10-16% difference between the actual running cost and the estimated cost given using this method for each of the appliance tested.

The work of Farinaccio & Zmeureanu, (1999) was to develop generic rules for two relatively high-energy consumption devices. The refrigerator would cycle on and off through the day and this would be affected by, for example, ambient air temperature and the number of times the door was opened. The water heater would be likely to have shorter periods of very high usage and would be switched on when the resident drew hot water. The experiment was limited in that it was only undertaken in one house with one set of appliances and the issue of seasonal variation was not considered in the experiment design.

A fourth example of appliance recognition is the RECAP (Recognition of Electrical Appliance and Profiling in Real-time) system developed by Ruzzelli, Nicolas, Schoofs, & O'Hare, (2010). This system aimed to overcome the issue of installation by using a single electricity monitor on the main electricity feed into the house. The electricity monitor used for this system was Episensor ZEM-30,[1] which can be used to capture information such as the real power, power factor, RMS (root mean squared) current, RMS voltage, peak (i.e., maximum) current and peak voltage.

The RECAP system was based on recording the signatures of individual appliances. This was undertaken by saving the data for each appliance, given by the Episensor, as well as the shape of the signal given by the appliance. These recorded signatures were then saved in a database. This database of signatures was used to train an artificial neural network. Once the neural network had been trained, it was used to recognise when the corresponding appliances were being used. Several different experiments were run using the RECAP system. The first involved recognising three kitchen appliances over one week (the appliance signatures were stored in the database before the experiment was carried out). This experiment gave an overall accuracy of over 95% for recognising when the three appliances being tested were turned on. Other experiments were undertaken to review the effectiveness of the system to differentiate between two appliances with similar signatures; this experiment was undertaken finding an appliance with a similar signature and introducing it into the same scenario as the first experiment. The

[1] http://episensor.com/products/

experiment was run over an hour and the system had an accuracy of over 84% for recognising when the three kitchen appliances were turned on.

The work of Lee, Lin, Jih, & Hsu, (2010) aimed to develop a model to recognise activities based on appliances and this information was to be used to highlight to the user where energy saving costs could be made. This work developed an appliance activity model framework, which could be used to associate the user's appliance usage to common activities. From analysis of the results from this model, the aim was to highlight to the user, appliances that were not being used at the time, i.e., were unattended, and therefore were wasting electricity and money. To develop the appliance-activity model the researchers developed a questionnaire asking the user what appliances they might use when performing certain activities, for example studying or preparing meals. This information could then be used to highlight appliances, which were on, but were not being used based on the information provided.

The experiments were conducted within an experiment room to recognise a list of nine appliances. Data were collected using a monitor, which collected the total electricity consumption of the experiment room every 5 seconds (Lin, Lee, Hsu, & Jih, 2010). For this, a Bayesian network was trained and tested to recognise the usage of these nine appliances. From this training and testing, the method showed a high rate of accuracy (>92%) for all the appliances. This model also recorded a high rate of precision (>83%) and recall (>76%) for all the appliances. For this work the researchers (Lee et al., 2010) only focussed on recognising appliances, with the aim of extending this method to highlight when the appliances were being left unattended in future work; however, this research collected information about what appliances people considered they might use when performing an activity. From the results of this data collection, the researchers highlighted that, on average, respondents considered that 23 'appliances' were used per activity, although some of these 'appliances' would be classed as objects, for example, a keyboard and a pen disk. (These examples were given as answers to the question of what appliances might you use when using a computer). The researchers further analysed the results given by these data and identified that only one appliance was used across all the activities that they surveyed, which was the lights. From these analyses the authors also identified that the majority of the appliances were each only linked to one activity.

Although this research does highlight that the use of appliances is activity related, the results from the survey are given in a very general context. This does not make allowances for the different appliances and the ways individuals may use different appliances in their own life.

A study by Lines, Bagnall, Caiger-Smith, & Anderson, (2011) collected electricity data from 187 houses for 12 months, using both individual appliance monitors and whole house consumption monitors, as described in section 2.4.1.3. The focus of their work was to develop a classifier to allow for automatic classification of a new appliance without any prior knowledge of that appliance. To address this, the authors employed a time series classification approach to classify appliances from their daily and weekly profiles. This approached involved collecting the data on appliances at 15-minute intervals. From each set of collected data, certain features were derived from the appliance profile, for example the mean, minimum and maximum values, and the standard deviation. The researchers evaluated a number of different classifiers, using both raw data and derived data from weekly and daily datasets. The results from a range of different classifiers showed that the derived features gave a better classification accuracy than the raw data, with the Random Forest classifier providing the best accuracy of 61.34% from the derived features and 59.04% from the raw data.

The work by Lines et al., (2011) highlighted that, although this method did produce good recognition from some appliances, for example, the kettle, other appliances were far harder to classify across the data, for example, the television and computer. To improve the accuracy, the researchers proposed combining certain appliances into similar groups, for example, combining the television and computer into an appliance group. This method did improve the accuracy of the different classifiers, with the best accuracy given by the Random Forest of 80.32% from the derived features and 72.96% from the raw data. However, the change also reduces the classifier’s ability to recognise some individual appliances.

The confusion matrix, which is discussed in greater detail in section 2.4.5.11, provided from the results from the work of Lines et al., (2011) also showed some confusion between the oven, the dishwasher and the washing machine. This confusion was not considered in the paper by Lines et al., (2011), although the reasons for the confusion between these appliances could lead to problems with the transferability of recognising appliances from multiple houses as well as the recognition of appliances with varying signatures. These issues are discussed in more detail in section 2.5.

2.4.3.3. Summary

This section has discussed the many different approaches, which have been used to collect and analysis electricity data so as to effectively recognise the use of appliances. From these approaches, there are different ways of collecting electricity data, from measuring multiple values, such as the electric current, as well as the power (Hart, 1992; Ruzzelli et al., 2010) or measuring single values, such as the total power consumption of a whole house (Lee et al., 2010; Lin et al., 2010) as well as from multiple appliances (Lines et al., 2011). The differences in approaches of data collection as well the method are discussed in more detail in section 2.4.7.

2.4.4. Activity recognition

The use of electricity monitoring devices not only allows the capture of the amount and cost of electricity usage, but also the data can show the time and duration of appliance use. This can be used to support the building up of a picture of a resident's typical use or pattern of living. Variations in the time or duration of use, for example in the case of sudden accident or illness, and may then be used to trigger an alert (and potentially an appropriate intervention).

The work carried out by Noury, Berenguer, Teyssier, Bouzid, & Giordani, (2011) aimed to monitor the trends of activity of elderly people living alone. To achieve this, an electrical monitoring device was placed into the electricity meter that was able to monitor the power consumption of individual devices and also when a device was turned on (it required a short learning period) (Berenguer, Giordani, Giraud-By, & Noury, 2008). To monitor the resident's activities, the project used activities of daily living (ADLs) (as discussed in section 2.3.6) but built and indexed these activities based on the electrical appliances used to carry them out (for example, switching the bathroom light on indicated going to the toilet or taking a bath) (Noury, Quach, et al., 2011). Two experiments were undertaken. The first was in 13 people's homes and the second in 12 people's homes. For the experiments, three ADLs (food preparation and eating, hygiene and toilet usage) were reviewed. These experiments also aimed to capture aspects of the resident's overall activities throughout the day and at night. The results from the experiments showed that an electricity monitor was an effective method that could be used to monitor resident's activities. Additionally, the monitoring of certain activities highlighted changes that were caused by changes in the resident's health (Noury, Berenguer, et al., 2011).

Franco, Gallay, Berenguer, Mourrain, & Couturier, (2008) also examined the use of electricity monitors to provide a non-intrusive way of monitoring resident's activities in their own homes. The project recorded the electricity usage of key located light fittings as well as key appliances (for example the toilet light, cooker) by placing electricity sensors on these items. The recorded data were then transmitted to an external server where analysis of the data was undertaken. An experiment was conducted involving 13 elderly people, one of whom had moderate Alzheimer's disease. For this experiment an average of 20 sensors were placed into each home and used to record activities and the rooms that the resident was occupying. The categorisations of the activities performed were broken down into specific time frames in order to differentiate between day and night activities. The participants were also required to complete a diary to record their activities (ADLs) throughout the day.

The results from this experiment showed that the electricity data recorded correlated well with the activities (ADLs) that the resident recorded in the diary. This experiment did, however, identify that monitoring certain items, such as low wattage lights (i.e., <40 Watts), was not very effective as the electrical monitors were unable to register the switching on or off of these items accurately. The experiment also showed that residents did not always turn lights on if they were only in a room for a short period (for example, not turning the light on to use the toilet); this lead to certain ADLs not being identified accurately.

2.4.5. Data mining, machine learning and pattern matching

2.4.5.1. Introduction

As highlighted in this literature review the automatic recognition of activities from data (e.g. from sensors, see section 2.3.7) has been achieved previously by using different machine learning or data mining techniques. This section will provide a description of these techniques and their approaches as well as providing some examples of their methods. This section will conclude with a discussion of some of the examples of these techniques that have been used in other similar research.

2.4.5.2. Data mining

Large amounts of data are being produced in the world today (Witten, Frank, & Hall, 2011) with the world turning electronic, the decreasing costs of disk and online storage, the easy access to online data and the vast amounts of data being collected and stored in a wide variety of fields, e.g., business, health, etc. (Kantardzic, 2011; Larose, 2005; Witten et al., 2011). Examples of some of these data include details of what products we buy in the supermarket (Witten et al., 2011) or our telephone usage (Hand, Mannila, & Smyth, 2001). With the growth of collecting data there is a need to analyse the data effectively to extract information and knowledge that could be of value (Hand et al., 2001; Kantardzic, 2011; Witten et al., 2011). An example of this is targeted marketing at a specific group of customers (Witten et al., 2011). The area of the process around the extraction of information from data is called data mining.

Data mining, as defined by Kantardzic (2011, p.6), *"is a process of discovering various models, summaries, and derived values from a given collection of data".* Kantardzic, (2011) and Tan, Steinbach, & Kumar (2006) highlighted that within data mining there are, generally, two primary goals, prediction or description.

Prediction data mining involves the use of the variables within the data to predict future variables. The aim of prediction data mining is to produce a model, based on the data, which can be used to perform further tasks such as classification or prediction (Kantardzic, 2011).

Descriptive data mining looks at finding patterns or relationships within the variables of the data, which can be used to describe the data. The aim of description data mining is to gain an understanding of the data through the patterns or relationships, which have been uncovered (Kantardzic, 2011).

Fayyad, Piatetsky-Shapiro, & Smyth (1996b) showed data mining to be a step in the overall knowledge discovery in database (KDD) process. They described the KDD process in a number of stages. These are:

1. **Data understanding**: Creating an understanding of the overall goals of this process as well as relevant prior knowledge of the data and the domain of the data (Fayyad, Piatetsky-Shapiro, & Smyth, 1996a; Fayyad et al., 1996b).

2. **Data selection**: This involves selecting a data set that will be used in the discovery process (Fayyad et al., 1996a, 1996b).
3. **Data pre-processing**: This involves cleaning the data to account for noise within the data, missing values within the data or outliers (Adriaans & Zantinge, 1996; Fayyad et al., 1996a, 1996b; Kantardzic, 2011)
4. **Data transformation**: This involves identifying a set of features from the data, which can be used to describe the data (Piramuthu, 2004). The feature set (as described in section 3.4.2) is chosen depending on what is the outcome of the overall goals of this process (Fayyad et al., 1996a, 1996b).
5. **Choosing a data mining method**: This involves the matching of the overall goals of the KKD process to a specific data mining method. Examples of some of these data mining methods are classification, clustering and regression (Fayyad et al., 1996a, 1996b; Kantardzic, 2011; Witten et al., 2011).
6. **Choosing a data mining algorithm**: Within the different data mining methods, for example, classification (Fayyad et al., 1996a, 1996b; Kantardzic, 2011; Witten et al., 2011), there are a number of different algorithms which can be used to perform the classification task, such as decision trees (Apté & Weiss, 1997; Witten et al., 2011). This step involves choosing a data mining algorithm that is appropriate for the data and the overall goals of the process (Fayyad et al., 1996a, 1996b).
7. **Data mining**: This involves applying the chosen data mining algorithm to the data.
8. **Interpretation**: This involves evaluating the results from the data mining step and visualisation of the results so that the user can understand the meaning in the results (Apté, 1997). This can also mean a return to any of the previous steps if the results from the data mining are not acceptable (Fayyad et al., 1996a, 1996b).
9. **Consolidation**: This step involves the implementation of the results from the process into another system (Fayyad et al., 1996a, 1996b) or a presentation and interpretation of the results (Kantardzic, 2011).

As the review of Kurgan & Musilek, (2006) highlighted, since the proposed structure of the knowledge discovery in database process by Fayyad et al., (1996a, 1996b) there have been a number of different KDD processes, which have been developed by researchers and industries.

2.4.5.3. Machine learning

Machine learning was described by Rogers & Girolami, (2012, p.1) as "*learning or inferring a functional relationship between a set of attributes variables and associated responses or target variables so that we can predict the responses for any set of attributes*". Therefore machine learning can be described as programming computers which, through learning, will automatically improve their accuracy using experience, i.e., example data or past experiences (Alpaydin, 2010; Mitchell, 1997).

Machine learning can be used to provide a prediction on future data points (Alpaydin, 2010; Murphy, 2012). The predictive type of machine learning is a supervised learning approached, which is discussed in more detail in section 2.4.5.6. Machine learning can also be used to provide a description of the data (Alpaydin, 2010; Murphy, 2012). The use of descriptive machine learning is classed as an unsupervised learning approach, which is discussed in more detail in section 2.4.5.7. As well as providing prediction of data and descriptions of data, machine learning can also be used to provide reinforcement learning (Alpaydin, 2010; Murphy, 2012). An example of reinforcement learning is training a computer to play a game and each time the computer wins it is positively rewarded and when it loses it is negatively rewarded (Mitchell, 1997).

Machine learning algorithms are used in a number of different areas and for a number of different tasks. Examples of these areas include pattern recognition (Theodoridis & Koutroumbas, 2009), classification (Alpaydin, 2010) and clustering (Murphy, 2012). Machine learning algorithms are also used within the process of data mining (as described in section 2.4.5.2), with the application of different machine learning algorithms to large datasets (Alpaydin, 2010).

2.4.5.4. Pattern matching

Pattern matching algorithms are use to 'match', either exactly or to represent similarity, to a user defined pattern across a data set (Sheik, Aggarwal, Poddar, Balakrishnan, & Sekar, 2004; Wang, Seidel, & Weinkauf, 2016). Pattern matching has been widely used on textual string data for example, DNA sequence matching (Chen, Lu, & Ram, 2004) or string matching for virus detection (Dang, Le, & Le, 2016), as well as other forms of data for examples vector data types (Ebling & Scheuermann, 2003).

2.4.5.5. Supervised and unsupervised learning

As highlighted in sections 2.4.5.2 and 2.4.5.3, different machine learning and data mining approaches can be used either to provide predictions on new data, or to provide description of the data. This has led to two different types of learning approaches, supervised and unsupervised learning.

For supervised learning, a training data set is provided, in which this training set consists of a set of input data with the corresponding correct responses (or targets) (Marsland, 2009).The aim of the supervised learning algorithm is to produce a model that can then be used to predict the targets of a new data set based on the input data and targets that the model has been trained to recognise (Kantardzic, 2011). An example of supervised learning method is classification, in which an input is placed into a class, based on the training and target data for each class (Marsland, 2009).

In contrast to supervised learning, for unsupervised learning, no target data are provided (Marsland, 2009). The aim of an unsupervised learning algorithm is to discover similarities (Marsland, 2009) or structures (Kantardzic, 2011) within the input data. An example of an unsupervised learning method is clustering (Theodoridis & Koutroumbas, 2009), in which the input data are grouped or clustered together based on their similarities (Marsland, 2009).

As well as supervised and unsupervised learning, there is also a third approach to learning, semi-supervised learning. For semi-supervised learning the inputs are provided with both target data and without target data (Witten et al., 2011). Semi-supervised learning is useful in cases of limited data with target data (Theodoridis & Koutroumbas, 2009; Witten et al., 2011), for example, where the small amount of data with targets can be used to classify and provide target data for the data without targets (Witten et al., 2011).

2.4.5.6. Examples of supervised learning methods

As discussed in section 2.4.5.5, for supervised learning the training data are provided with a set of their corresponding targets or labels (Witten et al., 2011), with the aim of producing a model based on the relationship between the training data and corresponding targets (Kantardzic, 2011). There are a number of different supervised learning methods that have been utilised, and this section will provide some examples of these methods.

2.4.5.6.1. Decision trees

Decision trees are a method of supervised learning which is used within data mining (Witten et al., 2011). Decision trees can also be expressed as classification rules or associated rules (Witten et al., 2011).

Decision trees are constructed using a logical method (Kantardzic, 2011), based on seeking the best attribute split that can be used to separate each of the classes provided in the training data (Witten et al., 2011), and this then continues on all subsequent splits until no further data splits are required (Witten et al., 2011). The structure of a decision tree consists of nodes, branches and leaves. Where each node is used to test a particular attribute (Witten et al., 2011), each branch attached to that node represents the possible outcomes from the node (Kantardzic, 2011), for example yes or no and finally, the leaves represent the classification. An example of a structure of a decision tree is shown in figure 2.1, where each circle represents a node, each line represents a branch and each square represents a leaf.

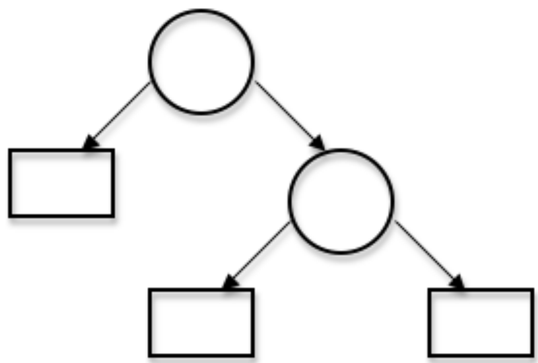

Figure 2.1: Example of a simple decision tree (Adapted from Murphy, (2012) and Witten et al., (2011))

Decision trees work well on simple problems, and have an advantage of fast computation time (Jain, Duin, & Mao, 2000) and on being easily interpretable (Kotsiantis, Zaharakis, & Pintelas, 2006). For more complex problems decision trees suffer from losing their interpretability, due their large size (Kantardzic, 2011) and suffer from over-fitting (Kotsiantis, 2013), a characteristic which is described in section 2.4.5.8.

2.4.5.6.2. Artificial neural networks

Artificial neural networks are inspired by biology and the study of how the brain computes and performs tasks (Kantardzic, 2011; Mitchell, 1997). The overall structure of an artificial neural network is modelled on the structures of neurons inside the brain (Kantardzic, 2011), with an artificial neural network containing a number of interconnected artificial neurons. An artificial neuron, as shown in figure 2.2, is constructed of three parts, the inputs with weights, an adder and an activation function (Haykin, 1999; Kantardzic, 2011). The inputs to the neuron are multiplied by their corresponding weights and passed to the adder, which sums up all the inputs. The summed weights and inputs are then passed to the activation function, in which, if the value of the summed weights and inputs are higher than the threshold of the activation function, the neuron provides an output (Theodoridis, Pikrakis, Koutroumbas, & Cavouras, 2010).

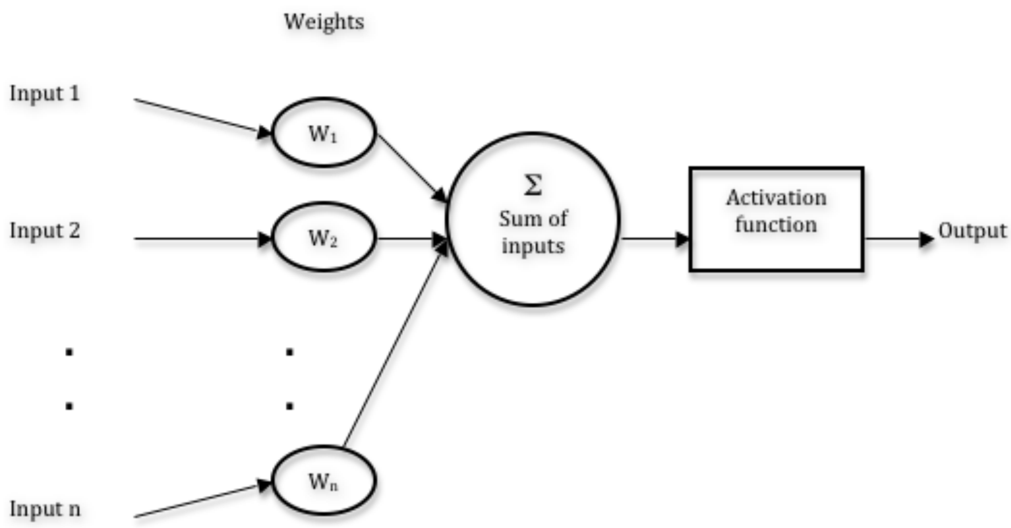

Figure 2.2: An example of an artificial neuron (Adapted from Haykin, (1999))

Artificial neural networks can be used for both supervised and unsupervised learning (Kantardzic, 2011). For a supervised learning task, a feed forward neural network is used, and this type of network is usually used for classification task (Gurney, 1997). A classification task involves placing input data into a class, based on the training and target data for each class (Marsland, 2009).

To form a feed-forward neural network, a number of the artificial neurons (as shown in figure 2.2) are joined, as shown in figure 2.3. Each circle, as shown in figure 2.3, is classed as node that contains the adder and activation from the artificial neurons and each connection represents the modified weights. The feed-forward neural network in figure 2.3 contains a three-layer structure, which consist of an input layer, a hidden layer and an output layer. The precise structure of a neural network is determined by the designer and, for example, can contain a number of hidden layers (Gurney, 1997).

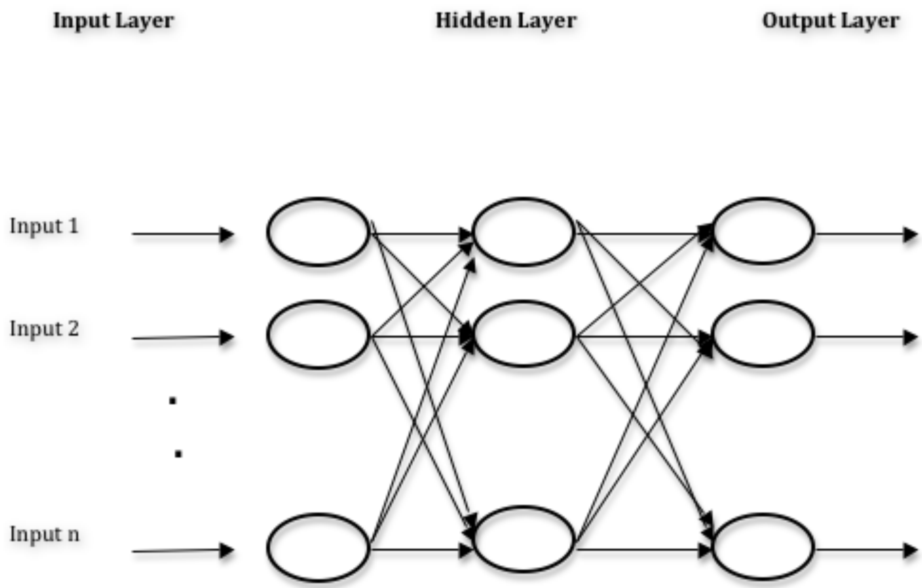

Figure 2.3: An example of a three-layer artificial neural network (Adapted from Haykin, (1999))

For supervised learning of a feed-forward network a back propagation algorithm is used (Duda, Hart, & Stork, 2001; Gurney, 1997). This process involves presenting an untrained neural network with the training data to provide an output. This output is then compared with the provided target data of the training data to produce an error (Duda et al., 2001), which then determines how the weights of the network are changed (Gurney, 1997). The process is repeated until the output of the model converge to match the provided target data closely (Duda et al., 2001; Gurney, 1997).

An advantage of using artificial neural networks is that they are good at finding non-linear solutions (Jain et al., 2000; Sathya & Abraham, 2013) as their structure allows the representation of non-linear decision boundaries (Witten et al., 2011). Similarly, to decision trees, as discussed in section 2.4.5.6.1, artificial neural networks also suffer from overfitting (see section 2.4.5.8).

2.4.5.6.3. Support vector machines

Support vector machines aim at finding the hyperplane which is a subspace (Cristianini & Shawe-Taylor, 2000) that provides optimal separation between two classes (Rogers & Girolami, 2012). This is achieved by finding the hyperplane which maximizes the margin, which is the distance between the hyperplane and the closest points on each side (Duda et al., 2001; Witten et al., 2011). For non-linear cases, the training of a support vector machine involves the transformation of the training data into a higher dimensional space (for example kernel functions (Rogers & Girolami, 2012)) where the data can be separated using a hyperplane (Duda et al., 2001). Support vector machines are used for binary classification (Rogers & Girolami, 2012), though they have been extended for multiple class classification tasks (Murphy, 2012).

Support vector machines have the advantage that they work well on small training datasets (Jain et al., 2000) as well as being less prone to overfitting (Jain et al., 2000) than artificial neural networks. However, with support vector machines, the risk of overfitting is increased with the addition of kernels (Cristianini & Shawe-Taylor, 2000). A disadvantage of support vector machines is that, compared with the other examples above, this approach does have a very slow training time (Jain et al., 2000).

2.4.5.7. Examples of unsupervised learning methods

As described in section 2.4.5.5, unsupervised learning methods involve the use of training data that does not have corresponding target data (Marsland, 2009). The aim of unsupervised learning is to find similarities in data (Kantardzic, 2011). Clustering is an example of an unsupervised learning approach, with the aim of clustering data into groups (Murphy, 2012). There are a number of different clustering algorithms (Duda et al., 2001), and one commonly-used type is k-mean clustering.

2.4.5.7.1. K-means clustering

The aim of the k-means clustering algorithm is to minimise the total square distances between the centre point and each of the other points in that cluster (Jain, 2010; Witten et al., 2011). There are a number of steps used for a k-means clustering algorithm. The initial step is to specify the ‘k’; the number of clusters into which the data points will be clustered. The second step is to assign a centre point randomly for each of the ‘k’ clusters. The data points are then assigned to their closest cluster (Jain, 2010), this being calculated based on their Euclidean metric between two points (Jain, 2010; Witten et al., 2011). The mean of each point in a cluster is then calculated, with this mean then become the new centre point for that cluster (Witten et al., 2011). This process is then iterated until there is no change in cluster assignment (Jain, 2010).

K-means clustering offers a simple (Witten et al., 2011) and easy to implement (Jain, 2010) method for clustering, although the value of ‘k’ needs to be specified and this might not possible in cases where the optimum number of clusters may not be known (Jain, 2010).

2.4.5.8. Overtraining and overfitting

As highlighted in section 2.4.5.6, many supervised learning algorithms such as feed-forward neural networks and decision trees suffer from overtraining, leading to overfitting. Overfitting is a term used to describe when a model has been over fitted to its training data (Rogers & Girolami, 2012) i.e., it produces a poor performance on new unseen data even though it produces a good result on the training data (Mitchell, 1997). There are a number of different methods that can be used to limit overfitting to training data, which are discussed in section 2.4.5.9.

2.4.5.9. Training and testing data

Training and test datasets are provided so as to train the model initially and then to provide an evaluation of their performance on an independent dataset, a test data set (Witten et al., 2011). There are a number different of methods of providing training and test dataset.

One method is the holdout method. This method splits the data set so as to provide a training set and a test set. A common split in data is 70% for training and 30% for testing (Witten et al., 2011). The advantage of the hold out method is that it provides training and test data if there are limited data for other training and testing schemes, such as a three-way split (described below). The use of holdout method has the disadvantage of not producing generalizable results, if either of the training or testing dataset are not representative of the overall data (Witten et al., 2011).

Another method for providing data for training and testing is to use a three-way data split. This involves splitting the data into a training, validation and test datasets. A validation set is a separate dataset, split from the original training set or provided separately, that is used to validate the predictive performance of the model (Rogers & Girolami, 2012). The use of a validation data set is used for overcoming overfitting of models (Rogers & Girolami, 2012) and the error produced from the test data set provides a good representation of the performance of the model on future datasets (Witten et al., 2011). The disadvantage of this method is that a larger amount of data is required to provide representative training, test and validation datasets and this may not always possible (Rogers & Girolami, 2012; Witten et al., 2011).

Cross-validation techniques are another effective method for reducing overfitting when data are limited (Rogers & Girolami, 2012; Witten et al., 2011). There are a number of different cross-validation techniques, with a common method being k-fold cross validation. This process involves the dividing of the dataset into a predefined 'k' number of folds, for examples three folds of data. The model is then trained on two of the three folds and tested on the remaining fold. This is then repeated so the model is trained on two of the three folds and tested on a fold that the model has not been tested on before. This is then repeated through all the folds until each fold has been used as test data only once. From each of the folds, the model produces an error rate based on the performance of the model on the test fold. The error rates from each of the folds are then averaged to provide a final error rate over all of the folds.

Another example of a cross validation method is the leave-one-out cross validation, this is similar to 'k' fold cross validation, although instead of k being a predefined number, in this method k equals the number of points in the dataset. Thus, the model is trained on all but one point of the data set and tested on the remaining point. This is then repeated on every point in the set. The overall error rate is given as an average of the error rate from each of the folds. This method has the advantage of training the model on the largest amount of data possible, and can be especially useful in cases of small datasets (Witten et al., 2011). A disadvantage of leave-one-out cross validation is that it is computationally expensive (Witten et al., 2011) and on large datasets it is not feasible (Rogers & Girolami, 2012).

2.4.5.10. Model evaluation

There are a number of different methods that can be used to evaluate the performance of a model based on its prediction on a test data set, either from cross validation or an independent sample.

2.4.5.11. Confusion matrix

For a classification task, a common method for evaluating the results of a model is to develop a confusion matrix of the results, as shown in figure 2.4.

<table>
<tr><td colspan="2" rowspan="2"></td><td colspan="2">Predicted class</td></tr>
<tr><td>Class 1</td><td>Class 2</td></tr>
<tr><td rowspan="2">Actual class</td><td>Class 1</td><td>TP</td><td>FN</td></tr>
<tr><td>Class 2</td><td>FP</td><td>TN</td></tr>
</table>

Figure 2.4: Example confusion matrix for binary classification

A confusion matrix provides an assessment of the results of a classier and in effect indicates how accurate the results produced by the classifier are (Kantardzic, 2011). A confusion matrix provides 4 measures, which are true positives (TP), false positives (FP), true negatives (TN) and false negatives (FN). Class 1 and class 2 are the two decisions to be made by a model (i.e., the prediction made by the model) and are compared with the actual class, i.e., the true value. True positives are the points when the classified has correctly identified a point. False positives are the points where the classifier has incorrectly identified a point. False negatives are the points where the classifier has incorrectly rejected a point. True negatives are points where the classifier has correctly rejected a point (Kantardzic, 2011; Rogers & Girolami, 2012; Witten et al., 2011).

From the values in the confusion matrix a number of calculations of the classifiers performance can be derived.

$$Accuracy = \frac{(TP + TN)}{(TP + FP + FN + TN)}$$

The accuracy of the classifier is the overall success of the classier to classify the data correctly, with the higher the value the more effective the classifier.

The error rate of the classifier can be calculated as 1- the accuracy. For some real world applications, the use of accuracy of a classifier can be misleading (Kantardzic, 2011). An example of this is in the case of disease detection (Rogers & Girolami, 2012), where those with the disease can be a small percentage of the dataset (1%) with the rest of the data set being healthy people (Kantardzic, 2011). In cases like this, the model can present a high accuracy without correctly classifying any of those with the disease. For these cases other measure can be used to indicate the performance of the classifier, such as the sensitivity (also called true positive rate) and specificity (also called false positive rate).

$$Sensitivity = \frac{TP}{(TP + FN)}$$

$$Specificity = \frac{TN}{(TN + FP)}$$

A perfect classifier would provide a sensitivity and a specificity both equalling 1, (although this is unrealistic). The use of sensitivity and specificity are a trade off between the two values which are dependent on the application of the classifier (Rogers & Girolami, 2012).

Other examples of values that can be used to indicate classifier performance, are positive predicted value (PPV) and negative predicted value (NPV). PPV represents the fraction of positive results, that are true positives (Murphy, 2012). With NPV representing the fraction of the negative results, that are actually negative results (Murphy, 2012).

$$PPV = \frac{TP}{(TP + FP)}$$

$$NPV = \frac{TN}{(TN + FN)}$$

Higher values of PPV and NPV can indicate a better performance of the classifier.

2.4.5.12. ROC curve

A receiver operating characteristic (ROC) is a plot of the true positive rate of a classifier (or sensitivity) against the false positive rate (or 1-specificity) for different values of the threshold of the classifier (Kantardzic, 2011; Rogers & Girolami, 2012).

A ROC curve is useful for evaluating the performance of a classifier using different learning methods (Kantardzic, 2011). From a ROC plot, the area under the curve (AUC) can be used as an indication of classifier performance (Rogers & Girolami, 2012), with the aim of finding a classifier with the highest value of AUC, (with a perfect classifier giving an AUC equalling 1 and a value of 0.5 or less being as good as random or worse).

A ROC curve and AUC provide good evaluation of the performance of a model as they take into account imbalances in data. A ROC curve is used for a binary or two case classification, therefore cannot be used to show performance for a multi-class classifier without representing the ROC curve as a series of one class against the remaining classes ROC plots. As highlighted by Rogers & Girolami (2012), showing the performance of multi class classifier using a series of ROC curves does provide useful information about the performance of each class, but it is not clear how the multiple AUC calculations could be used to represent an overall generalizable score for the classifier.

2.4.5.13. Machine learning and data mining for electricity data analysis

Within the areas of appliance recognition from electricity data, different machine learning and data mining techniques have been utilised. The aims of these techniques are to recognise automatically when an appliance has been used based on a set of electricity data collected from a household. Some examples of these different techniques have already been discussed in section 2.4.3.2; this section will discuss, in more detail, the specific use of data mining and machine learning techniques that have been utilised for identifying the use of appliances.

Ruzzelli et al., (2010) used an artificial neural network as the basis of their RECAP system (discussed in section 2.4.3.2). This work highlighted some of the positive and negative aspects of using an artificial neural network for the recognition of electrical appliances. Among the advantages of using an artificial neural network for this task were that an artificial neural network could handle different data types and that an artificial neural network has a framework that could be easily adapted, i.e., to add more inputs (Ruzzelli et al., 2010). For this task, one of the major disadvantages of the use of an artificial neural network was the time that it took to train the artificial neural network to recognise appliances (Ruzzelli et al., 2010). This was further exacerbated with appliances which themselves had long running times, for example a washing machine.

For the appliances that were trained to be recognised using this method, there was a 95% recognition rate for three ON/OFF appliances: the kettle, the microwave and the fridge. However, it was noted that these appliances had a signature much higher than every other appliance and the lights that were also activated during the same training and testing period (Ruzzelli et al., 2010). This highlighted that, although recognising individual appliances in isolation produced a good recognition rate, the same might not be true when similar appliance signatures were present in the data. To test the effect of having appliances with similar signatures a fourth appliance was included in the study, i.e., an electric heater, which had a similar signature to the kettle and the microwave. The addition of this fourth appliance reduced the recognition rate to 84% from 95% recognition rate for three appliances.

For this method, it was noted that this artificial neural network performed well with the recognition of high-powered ON/OFF appliances, although the tests were conducted within a test environment, not in a real environment. This model was not trained to recognise more complex power-changing appliances, for example an electric oven. This was commented on by the authors as an area of further work, as well as an area that could be more complex were it to be investigated (Ruzzelli et al., 2010).

In a different study, Spiegel & Albayrak, (2014) compared the results from a number of different machine learning techniques applied to the same dataset. The dataset, as described by Kolter & Johnson (2011), consisted of data from six households recorded from the two main phases and from each individual circuit at a frequency of 1Hz. The recording of this data required a much larger amount of equipment to be installed (Kolter & Johnson, 2011), than the single electricity monitoring installed in the study by Ruzzelli et al., (2010).

The four machine learning techniques utilised by Spiegel & Albayrak, (2014), were a Naïve Bayes classifier, a factorial Hidden Markov Model, a classification tree and a one-nearest neighbour classifier. A Naïve Bayes classifier, as described in more detail in section 3.4.4, is a probabilistic classifier based on Bayes rule and certain assumptions (Stone, 2013). A hidden Markov model is used for probabilistic modelling of time series data (Ghahramani & Jordan, 1997), which is the application of a Markov chain with hidden states (Spiegel & Albayrak, 2014). A classification tree is discussed in more detail in section 2.4.5.6.1. A one-nearest neighbour classifier, is a simple classifier (Murphy, 2012) and can be seen as a baseline method for classification (Spiegel & Albayrak, 2014). The process of one-nearest neighbour is to compare a new instance point to all existing instance points with a distance metric (Witten et al., 2011), the points, and therefore the class which the new instance is closest to, is the class to which the new instance is assigned (Witten et al., 2011).

Each of these methods had the characteristics of fast fitting speeds and medium to fast predictions speeds for large datasets (Spiegel & Albayrak, 2014). The comparison of the different methods highlighted that, overall, the Naïve Bayes classifier gave the best mean accuracy across all of the appliances, although the accuracy for each individual appliance did vary from 77.51% to 99.82% (Spiegel & Albayrak, 2014). For each of the households the Naïve Bayes also gave the best mean, with a value of 89.42%, although this varied between the households from 81.31% to 98.64%.

It was noted by the authors that each of the different machine learning methods performed well on appliances that are distinctive in their power change profiles, for example, the microwave or air conditioning (Spiegel & Albayrak, 2014). However, the different methods performed less well on appliances that had changes in their signatures that could be confused with other similar appliances, such as lighting or the refrigerator (Spiegel & Albayrak, 2014).

Chapter 3 of this thesis discusses in more detail, i.e., from a methodological perspective, the steps involved in developing a model used to recognise appliances from electricity data that was chosen for this research. Having described the range of data mining methods, the following section discusses the privacy concerns around using electricity monitors to monitor resident’s activities.

2.4.6. Privacy concerns of electricity monitoring

Within the work discussed in section 2.4.3, researchers did not highlight any privacy or ethical considerations in their work. However, for all of these examples, when data had been collected to conduct experiments, the data were collected from test environments and not from peoples’ actual homes.

Within the work discussed in section 2.4.4, the researchers did highlight some privacy and ethical considerations in their work. As highlighted by Franco et al., (2008), data security is an important issue in the collection of data from monitoring devices placed within peoples' homes. For their work, the data collected from the monitoring devices were encrypted and the researchers required passwords to access the data. Consent was also required from all of the participants and they were also able to withdraw their consent and leave the data collection at any time. The work of Noury, Berenguer, et al., (2011) also made similar privacy and ethical considerations in their collection of data from peoples' homes.

As shown by the research in sections 2.4.3 and 2.4.4, the extent to which residential electricity data can be analysed and used to highlight different aspects, such as appliance usage, activities and patterns is still being examined by researchers. The concerns as to what this data can show and what it can be used for were highlighted by Sintoni et al., (2011). This work discussed that, with prior knowledge of electricity appliances within a house, data that has been collected by electricity monitors could be used to show information about the user that they themselves might not be aware was being collected. This, as the research highlighted, could be used to learn the activity patterns of the occupants without their knowledge. However, for the example given in the research described by the work Sintoni et al., (2011), data had been collected and analysed with the full knowledge and consent of the occupants of the house.

The work of Kolter & Johnson, (2011) also highlighted some privacy concerns of the use of electricity data. Kolter & Johnson, (2011) collected the data via non-intrusive load monitoring (as described in section 2.4.3.1) from a number of different houses. The data collected by these researchers was made publically available, although the researchers did highlight some privacy concerns with the release of data of this kind to the public. They highlighted that sharing of a real time data set, as well as the location, could be very harmful to the occupants as it would be possible that analysis of the data could give an indication of whether the occupant was at home or not, based only on their current electricity usage. To address these concerns the researchers decided on a number of safeguards to protect the privacy and safety of occupants. The first of these safeguards was that no information was stored about the location of the houses and no information about the houses was released, only the city. The second safeguard was that only historical data were made publically available, with the real time data only available to the researchers. Although it was noted by the researcher that privacy concerns from this type of data does require constant monitoring, with the safeguards that the authors applied to this data they believed that the risk of revealing personal data or their location of the houses involved in this data collection was low.

2.4.7. Summary

The work discussed in section 2.4.1- 2.4.6 has shown the potential for using a single non-intrusive sensor as a means of not only monitoring resident's energy and appliance usage, but also the activities being undertaken. This work has also highlighted a range of potential difficulties when trying to monitor a resident's activity using electricity data. The work carried out by Hart, (1992) grouped electrical appliances used in people's homes into four groups. Only two out of the four groups of appliances can be effectively monitored; however, it would appear that most of the appliances linked to the identified activities fall into these the two recordable groups. The work by Hart, (1992) also highlighted that some appliances go through different states depending on the settings to which they are set. A good example of this is a washing machine that has a different setting depending on which type of clothes the resident is washing, e.g., wool versus cotton. The ability to monitor appliances that have different states have had varying accuracy, depending on the appliance and the number of settings each has. Importantly, all of the different approaches to both appliance recognition and activity recognition have required initial training of the system to learn underlying baseline information. Within this, both Hart, (1992) and Franco et al., (2008) have noted that it is very hard to achieve effective training of the systems to identify low energy devices accurately.

The work of both Franco et al., (2008) and Noury, Berenguer, et al.,(2011) outlined the monitoring of residents' activities using ADLs based solely on their electricity usage of different devices. This work has shown the possibility of using an electricity monitor as an additional sensor that could be used for monitoring. However, this method of monitoring required a large number of individual electricity sensors placed on a number of devices (in the case of Franco et al., (2008)) or required hardwiring of a device into the electricity meter of the house (in the case of Noury, Berenguer, et al.,(2011)).

The work by Farinaccio & Zmeureanu, (1999) and RECAP (Ruzzelli et al., 2010) showed that each electrical appliance has a different signature and that with the recording and learning of those signatures it is possible to break down electricity usage data into the appliances being used. However, to achieve this, prior knowledge or a training period is required before the method can recognise subsequent appliance usage.

Within the research, different types and numbers of sensors have been used to monitor different variables within electricity data though with the aim of recognising when appliances have been used. Examples of these differences are the work of Franco et al., (2008); Lee et al., (2010); Lines et al., (2011) and Ruzzelli et al., (2010). The work of Lee et al., (2010) used a single sensor to monitor a single variable (power consumption), whereas the work of Franco et al., (2008) installed multiple sensors on electrical appliances, to indicate when the appliances have been used. Finally, the work of Ruzzelli et al., (2010) used a single sensor to record multiple variables within the electricity data, such as current.

In summary, this section has reviewed work that has focused on providing the recognition of different appliances from a single non-intrusive electricity monitor. The section has also reviewed work on the monitoring of resident's activities (with an ADL scale) from electricity data provided by a number of electricity sensors linked to appliances throughout a house and a device hardwired into the mains electricity meter.

2.5. Synthesis and gaps in the literature

This chapter has reviewed the literature around the areas of health smart homes (section 2.3.3), telecare (section 2.3.4) and electricity data analysis (section 2.4). This review has helped in the highlighting of areas of discussion and gaps in the literature that provide the investigation for this thesis. This section will draw on the reviewed literature to provide a discussion and highlight the gaps in the research that are further investigated in this thesis.

2.5.1. Development of the list of activities

From the literature review it is shown that, to develop a robust process for using sensor technology to monitor activities, a systematic approach should be taken to develop or choose the list of activities, which the sensor(s) can be used to monitor. This section will discuss the approaches undertaken in previous research, highlight the gaps in the research and justify the approach taken for this research.

As shown in this literature review, previous research have adopted one of two approaches in choosing activities to be monitored, ADLs (as described in section 2.3.6) or developing a list of activities, as highlighted by the work of Glascock & Kutzik, (2006, 2007). As highlighted in section 2.3.4, Glascock & Kutzik, (2006, 2007) developed a system that did effectively identify changes in behaviour that could be attributed to changes in medical condition based on a list of general activities. In contrast, Alwan et al., (2006) highlighted the potential of using a variety of sensors to monitor changes in the performance of certain ADLs, which could be explained by changes in the resident's health. Both of these approaches have highlighted changes in the activities being performed that could then be explained by changes in the health of the resident. However, it is not clear from the research which, if either method provides a more robust way of highlighting any potential change in health of the resident. It could be argued that it is just as beneficial for the relative or caregiver to know the daily ADL score for the resident as it is to know that they had got out of bed at a reasonable time in the morning and made breakfast.

Therefore, in this research the approach in developing a lifestyle or behavioural monitoring system is to focus on recording activities that, if changes in frequency or time of these activities are shown to occur, the reasons for this change maybe reasonably associated with a health change of the resident (Brownsell et al., 2011). It will then be important to be able to present information or 'changes' in the resident's activities in a way that can be interpreted by a caregiver and also keeping false alerts to a minimum. With caregivers or relatives potentially being the main users of these types of system, the focus of developing the system and the activities chosen should concentrate on the areas that caregivers or relatives would identify as being important to be monitored.

As highlighted in section 2.3.8, within the areas of telecare and the use of telecare systems, the views of those who have the technology placed into their homes have been well documented, for example the work of Demiris et al., (2008), Milligan et al., (2011) and Sixsmith et al., (2007). There is, however, limited research, using case scenarios in the work of Percival & Hanson, (2006) into the views of carers and relatives into what features they may want from a remote monitoring system. To address this gap in the research, it was therefore decided to conduct a survey into the key areas and information that carers/relatives, who are likely to be key potential users of the information, would like to have access to in order to be reassured about their elderly/ill person. The results from the survey were then used to inform a list of activities to be monitored using a whole house electricity consumption monitor, thus developing a list of predefined activities rather than using ADLs. The reasons for this approach are:

1. To use ADLs it is necessary to score a range of measures, some of which are not related to electrical use and because the project is to be based solely on monitoring electricity it is not possible to measure all ADL activities. An example of this is getting dressed, which requires no electricity usage.
2. Measuring ADLs requires active and on going measurement and the concept of this work is to review non-intrusive lifestyle monitoring, and so ADL will not fit this concept.

As well as the development of a list of activities for monitoring, previous research, for example the work of Bowes et al., (2012), Demiris et al., (2008) and Stowe & Harding, (2010), have highlighted certain considerations that need to be made in terms of privacy and the ethics of monitoring information. These issues were highlighted in sections 2.3.8, 2.3.9 and 2.4.6 and should be reflected in the design and implementation of a monitoring system. The first of these is data security; the data that has been collected from the sensors must be stored and transmitted securely. Secondly, confidentiality must be assured, as the information recorded by the sensors must only be disclosed to those whom the user has agreed. Thirdly, informed consent must be in place, i.e., the user must consent to having a monitoring system placed within their home and must be free to withdraw and remove or turn off the monitoring equipment at any point. Finally, the goals of the monitoring system must be clearly defined, with the researcher and the user clear about what the sensors will monitor and what the data collected is used to monitor. How these steps are implemented in the design of the monitoring system for this thesis is discussed in Chapter 3.

This section has highlighted the gaps in the research into the development of the list of activities to be monitored; the next section will discuss the previous research using sensors and highlight the gaps in the research of using a single electricity sensor to monitor activities.

2.5.2. Monitoring of activities from a single whole house electricity monitor

The second point of discussion and investigation highlighted from the literature review is the use of a single whole house electricity monitor to monitor the activities of the residences within their home. This section will discuss the approaches undertaken in previous research, highlight the gaps in the research and justify the approach taken for this research.

As discussed in section 2.4.3 and 2.4.4, electricity data can be provided by a number of different sensors or by a single sensor and can be used to show appliance usage and activities. Recording electrical data from a range of multiple sensors placed on appliances across a house has the advantage of being able to record a large amount of information, for example individual appliance usage from each plug socket in the house (see section 2.4.1.1) as well as to capture information at very fast frequencies (Zeifman & Roth, 2011). However, the use of multiple sensors has the disadvantage of the cost of placing large amounts of equipment into the resident's homes. There is also an issue of intrusiveness, as highlighted by this review in section 2.3.8, the installation of equipment into homes can be intrusive to the user, and can lead to them not wanting to participate (Bowes et al., 2012). The collecting of electricity data for example in the works of Franco et al., (2008) and Lines et al., (2011) could be seen as intrusive as the approach of both these researchers required the installation of a large number of sensors to monitor each individual appliance. A way to address this issue is to design a monitoring system that is non-intrusive to the user, and does not require the users' input to operate it. The use of a single whole house electricity monitor, as a non-intrusive sensor would provide a compromise to this issue, as the monitor can be easily installed into a house. The installation of this single sensor is also low cost and can be almost invisible to the user as their electricity usage is recorded indirectly, with the electricity sensor placed around the mains electricity fuse box, which is usually out of sight.

As shown in the literature in sections 2.4.3 and 2.4.4, there are many different monitoring devices, as well as different types of information that each of these devices can capture. This makes it difficult to assess the different approaches to analysing electricity data as the data changes greatly between the different monitoring devices. For this thesis, the data will be collected by a single electricity consumption monitor, which will provide the power consumption of the household, recorded at 6-second intervals. The single whole house electricity monitor (as shown in section 3.3.2) used for this thesis is similar in design to that used for the data collection in the work of Ruzzelli et al., (2010), although the data that is captured is different. The data which is captured will be similar to that captured by the work of Lee et al., (2010), although the monitor used is different in the design and thus the monitor used for this research is easier to install.

From the work of previous researchers as outlined in section 2.4.3 and 2.4.4, the collection of whole house electricity consumption data from only a single monitoring device has only been collected from single houses or from test environments, and not in multiple houses in the real life situations. This is a gap in the research that this thesis will aim to address by collecting information from multiple houses; this study is described in Chapters 5 and 6 of the thesis.

Within this study, the aim is to also investigate the <u>feasibility</u> of collecting electricity data from multiple houses, as well as to consider the <u>transferability</u> across multiple households. Lee et al., (2010) and Ruzzelli et al., (2010) both showed good recognition rates using one set of appliances, which their systems were trained to recognise. However, neither of these studies addresses the transferability of their systems across data collected from multiple households, different appliances or the same appliances from different manufactures. This is a limitation of these works as the transferability of these systems were not assessed and is key to the wider development of this approach. The work of Lines et al., (2011) use multiple electricity consumption sensors from multiple houses to recognise appliance usage across 187 houses although, as highlighted in their results some appliance types, for example washing machine or oven, produced less reliable results than other appliances. Lines et al., (2011) did not comment on the less reliable results for some of the appliances across the households, though the reasons for these less reliable results could highlight a lack of transferability of the classifier across multiple households with the same types of appliances. This study will investigate the transferability, of the classifier to analyse data from multiple households with both the same appliance type as well as different appliances, as described in Chapters 5 and 6 of the thesis.

The challenges of this type of monitoring, which have been highlighted from this review, will also need to be considered within the design and implementation. As shown in section 2.4.3 and discussed by the work of Hart, (1992), appliances can be placed into four categories, with only two categories being able to be effectively recognised. This review has highlighted the previously documented problem (Franco et al., 2008) of recognising low power appliances. This needs to be considered when choosing a list of appliances to be monitored. In addition, the use of gas appliances for some activities needs to be addressed, as depending on the house or their habits some activities might not use major electrical appliances (such as the oven). This could therefore limit the effectiveness of this approach to monitoring and will need to be addressed in this thesis.

This section has highlighted the gaps in the research into using a single whole house electricity monitor to monitor the activities of the resident. Chapter 3, Chapter 5 and Chapter 6 of this thesis, will discuss the steps taken to use a whole house electricity monitor (as discussed in section 2.4.1.2) as a non-intrusive remote monitoring sensor, to monitor the user's activities based on their appliance usage. Chapter 5 and Chapter 6 of this thesis will also discuss the issues of transferability, of a classifier, across multiple households as well as the effect different appliance signatures for the same appliance (from different manufactures) has on the recognition of the appliance.

The next chapter of this thesis (Chapter 3) will go on to discuss the methodology and methods that were utilised for this research.

Chapter 3: Methodology

3.1. Introduction

Chapter 2 of this thesis highlighted and discussed the previous research that has been undertaken in this area. The conclusion of this literature review (section 2.5) highlighted some of the gaps in this research and discussed how these will be addressed by this thesis. This chapter will follow on from the conclusion of Chapter 2 and aims to provide an overview of the methodology used for this research and also the different methods used to collect the data. As described by Blaxter, Hughes, & Tight, (2010) a methodology refers to the overall research approach and incorporates the theories behind the research, as well as the research methods. In contrast, a research method refers to the type of tools that are used to collect data for research, for example surveys or interviews (Walliman, 2011). Further details on the methods used within each part of the research are included in the subsequent chapters.

This chapter is divided into three parts. Section 3.2 gives a short general overview of the different research methodologies and a description of the methodology used for this research. Section 3.3 provides a description of the different data collection methods used for this project. Section 3.4 discusses some of the different issues that need to be considered when choosing a technique to analysis electricity consumption data. Section 3.4 also gives a description of the method used to analyse the electricity consumption data for this research and how this was implemented. As mentioned above, further details of the methods used for the data collections are provided in their relevant chapters.

3.2. Research methodology

This section gives an outline of the methodological approach used in this research. Further detail of each of the methods will be provided in the subsequent analysis chapters, 4, 5 and 6.

3.2.1. Research philosophy

Research philosophy influences the practice of research (Creswell, 2014). How research is conducted is deeply influenced by the philosophy which is used to underpin it (Walliman, 2011). There are four main philosophy approaches or paradigms (Creswell, 2014) that are discussed in more detail in sections below.

3.2.1.1. Research paradigm

The term research paradigm, as defined by Bryman, (1988, p.4) is a *"cluster of beliefs and dictates which for scientists in a particular discipline influence what should be studied, how research should be done, how results should be interpreted*". As stated simply by Punch, (2005, p.27), "*it means a view of how science should be done*". The four paradigms which will discussed further are positivism, post-positivism, critical theory and interpretivism.

3.2.1.2. Positivism

Positivism as described by Barron, (2006, p.212-213) "*advocates the application of the methods of the natural science to the study of social reality*". Therefore, the positivist view argues that scientific measures can be utilised to measure human behaviour similarly to those utilised to measure natural science (McNeill & Chapman, 2005). For positivists, knowledge is derived from scientific methods (Walliman, 2011). From this, the knowledge gained can be used to build cumulative parts which add to what is already known (Walliman, 2006, 2011). The different methods utilised for conducting positivist research are closely related to methods utilised for conducting quantitative research (Punch, 2005). Some examples of the different methods are surveys and experiments (Barron, 2006).

3.2.1.3. Interpretivism

Interpretivism, is the contrasting paradigm to positivism (Bryman, 2012). The interpretivists share the view, as discussed by Bryman, (2012, p.28), "*that the subject matter of the social sciences - people and their institutions - is fundamentally different from that of the natural sciences*". Therefore, the subjective experiences (McNeill & Chapman, 2005) and meanings are critical to social actions (Walliman, 2006). From this, the aim of interpretivism, is to understand the world as their research subjects do (McNeill & Chapman, 2005) and using this to draw interpretations and meanings (Walliman, 2006). The methods used to conduct interpretivism research are those typically used for qualitative research (Barron, 2006). Some examples of these different methods are participant observations or unstructured/semi-structured interviews (Barron, 2006).

3.2.1.4. Post-positivism

Post-positivism has emerged as a reaction to criticism of the positivism paradigm (Creswell, 2014). Post-positivists hold the view, as described by Sharma, (2010, p.702), "*that humans are biased in their perceptions of reality and that hence we can approach the truth of reality but never explain it fully*". Therefore, in the views of post positivists, the absolute truth can never be found (Creswell, 2014; Sharma, 2010) which is in contrast to the views of positivists. The research conducted by post-positivist focuses on examining the causes that influence outcomes (Creswell, 2014). Within paradigm, different research methods are utilised in combination (Sharma, 2010), for example, using both quantitative and qualitative research methods (Pickard, 2013).

3.2.1.5. Critical theory

Critical theory as described by Howell, (2013, p.81), "*involved ideas relating to the empowerment of the people; it should challenge injustices in social relations and social existence*". Some of the examples of social injustices that are challenged are racism, gender inequality and class inequality (Creswell, 2014). Critical theory utilises both qualitative and quantitative research methods as well as mixed methods (Willmott, 2008), as described in section 3.2.3.

3.2.1.6. Summary

This section has provided a short overview of the different research paradigms and research methods for each used within social science research. The research in this thesis follows the post-positivism paradigm, as the use of quantitative and mix method approaches are traditionally aligned with post-positivist paradigm. The data use in this thesis, from the survey and analysis of electricity consumption data are quantitative and thus aligns with the post-positivist approach. Section 3.2 and 3.3 providing an overview of the different methods used for this research.

3.2.2. Qualitative and quantitative research

Within research, two of the most common research methodologies are qualitative research and quantitative research. The distinction between these two methodologies are highlighted by Flick, Von Kardorff, & Steinke, (2004) and McQueen & Knussen, (2002). Flick et al., (2004, p.3) described qualitative research as a way to "*describe life worlds 'from the inside out' from the point of view of the people who participate. By so doing it seeks to contribute to a better understanding of social realities".* Whereas the work of McQueen & Knussen, (2002, p.27) says that "*Quantitative research reflects the philosophy that everything in the social world can be described according to some kind of numerical system*". A simplistic distinction between these two areas of research is highlighted by the work of Punch, (2005, p.3) where *"Quantitative research is empirical research where the data are in the form of numbers. Qualitative research is empirical research where the data are not in the form of numbers"*.

3.2.2.1. Quantitative research

As highlighted by Punch, (2005), McQueen & Knussen, (2002) and Walliman, (2011), quantitative research aims to gain information from numerical data by using different types of numerical analysis, for example descriptive and inferential statistical analysis (Foster, Diamond, & Jefferies, 2015). The data used for quantitative analysis does not have to be numerical in its structure (for example the number of hours of TV watched each week) but can be used to represent a fixed group of responses, for example gender (McQueen & Knussen, 2002; Walliman, 2011).

The advantages of using quantitative research methods are that they help to provide answers for the "what" type of research questions as they can provide good descriptive results (Patten, 2007). Quantitative research is generally easy to reproduce (Bryman, 2012) and to generalise the results to the wider population (Patten, 2007; Walliman, 2011) , assuming the sample is representative of the wider population. The disadvantages of using quantitative research methods are that they are not particularly good at answering "how" or "why" research questions (Blaxter et al., 2010; Patten, 2007) and the response rates need to be high so that the sample can be considered representative of the population (Patten, 2007). In addition, minimum sample sizes are required for statistical tests to be valid (Bryman & Cramer, 2011; Walliman, 2011; Yates, 2004).

3.2.2.2. Qualitative research

Qualitative research aims at gaining an understanding into different behaviours and the possible reasons for them (Flick, 2014); examples include people's emotions, ideas, fears and beliefs. The data used in qualitative research are generally in the form of words (Patten, 2007) and therefore cannot be counted to form mathematical measures or statistics. Examples of qualitative research methods include using interviews and focus groups (Blaxter et al., 2010; Flick, 2014; Walliman, 2011).

The advantages of using qualitative research methods are that they help to provide answers for the "how" or "why" research questions (Patten, 2007) and they allow researchers to explore more into the reasons behind the answers (Flick, 2014). The disadvantages of qualitative research are that the findings of the research can be subjective (Blaxter et al., 2010) and open to different interpretations and so qualitative research does not seek to generalise the results across a population but to produce findings that can be transferred across groups in similar situations (Blaxter et al., 2010; Patten, 2007).

3.2.3. Mixed methods research

A mixed methods approach to research is the combinations of two or more different research methods within one research project (Blaxter et al., 2010). Generally, a mixed methods approach uses combination of qualitative and quantitative research methods (Blaxter et al., 2010; Bryman, 2012), although it can use different quantitative or qualitative methods.

3.2.4. Approach adopted in this research

This research used a mixed methods approach of combining a quantitative research method, i.e., a web-based survey, with the inclusion of some questions that could be analysed qualitatively, followed by the collection of whole house electricity consumption data and appliance usage diaries (as described in sections 3.3). The next section of the chapter will discuss the different methods used to collect the data used for this project.

3.3. Data collection: survey and electricity

This research has three data elements, a web-based survey, and the collection of electricity data combined with recording the use of electrical appliances in a diary. Section 3.3.1 gives an overview of survey theory and the advantages and disadvantages of using surveys for data collection. Section 3.3.2 describes how, as part of the project, data were collected from a number of households, over a one-week period to allow the recognition algorithm (allowing the identification of different electrical appliances from electricity consumption data) to be developed, trained and refined.

3.3.1. Survey theory

The advantages of using a survey for research are that they are quick, cheap and easy to distribute to a large number of people (McNeill & Chapman, 2005), particularly if using a web-based survey. The disadvantages of using a survey for research are that a representative sample and a large response rate are needed to be able to generalise the results (Patten, 2007). In addition, because surveys have a rigid format, there is no way to gather information about why people have given a particular response (Oppenheim, 2000). There are also a limited number of questions that can be asked in a survey, as a large survey can be off-putting to the respondents (Bryman, 2012) and can therefore limit the response rate. The obvious limitation of a web-based survey is that the respondents need to have access to the Internet to respond, so may lead to non-response bias.

A survey can be designed in a number of ways with various types of question structures (Oppenheim, 2000). The survey used for the first phase of this research contained both closed and open questions and also scaled and ranked question types. Closed questions contain a limited number of responses, e.g., with “yes” or “no” answers. Closed questions also include multiple-choice questions, which ask the respondents to choose the appropriate answer from a list (Walliman, 2011). Open-ended questions provide space for free text answers for the respondent to answer the question as they wish (Walliman, 2011). The third type of question is a scaled question, which asks the respondent to indicate how much they agree with a statement (Bryman, 2012). The final type of question used in this survey is a ranked question, which asks the respondents to compare a list of statements and rank them as they see appropriate (Bryman, 2012; Walliman, 2011).

As highlighted in the literature reviewed in section 2.5, previous research has adopted two approaches for analysing the data provided from sensors; these are the use of activities of daily living or a predefined list of general activities. Previous work by both Glascock & Kutzik, (2006, 2007) and Fleury et al., (2010) have provided examples of the two different approaches of analysing sensor data to provide an ADL score or analysing a list of predefined activities to provide a summary of activities. However, a revised approach to analysing sensor data could be to capture what would be seen as a normal pattern of usage and then to analyse the data to identify deviations from this normal pattern. However, capturing the normal electricity consumption usage could be very complex due to the large range of variability in the electricity usage. Therefore, instead of looking at deviations of the overall electricity usage, the recognition of individual appliance could be used to show the performance of specific, or more general, activities and thus provide a carer with assurance that the activity has occurred. From this, the data could then be used to highlight a deviation in a normal appliance usage pattern, for example, if a person habitually uses the oven to cook their meal every evening, and then they suddenly stop using the oven for a number of days, it might indicate a change in behaviour, possibly due to a sudden health problem. Thus, a pattern of the use of individual appliances could be monitored and analysed, with a view to recognising individual activities, rather than an overall pattern of electricity usage.

Therefore, and as is discussed in more detail in Chapter 4, a survey was designed to collect the views of relatives and carer into what activities they would like to know that their relative has undertaken. This was deemed to be preferable to the researcher selecting a list, which might not include activities considered to be important by carers, or using a pre-defined list of ADLs, which might not require electricity to be used, e.g., dressing, using the toilet. The intention was to use the results of the survey to inform a list of activities that could then be recognised from electricity usage, e.g., using the cooker or boiling a kettle, based on the activities that the relatives or carers would like to know that the person had undertaken, e.g., making a meal or having a drink.

Chapter 4 of this thesis discusses the method used to collect the survey data as well as the analysis and the interpretation of the results. It also considers how the results of the survey were used to identify which activities were deemed important to relative and carers, and therefore to inform the collection of the electricity data to identify appliances.

3.3.2. The electricity data

The second part of the data collection for this research involved the collection of two sets of data over a one-week period. These were the collection of whole house electricity consumption data and diaries of appliance usage. To allow for the collection of the electricity consumption data, appropriate equipment, as outlined in section 3.3.2.1, was placed within the participants' households. The method undertaken for the collection of the whole house electricity consumption data as well as the receiving and the storing of the data is also outlined in section 3.3.2.1. The method for the collection of the diaries of appliance usage is described in section 3.3.2.2.

The set up of this experiment had a number of limitations within the working conditions of the electricity monitor and data logger. These were:

1. The participants had to have access to their electricity meter, so as to be able to install the equipment.
2. There is an operating range for the electricity monitor (as highlighted in section 3.3.2.1.2) to receive the signal from the mains sensor (as described in section 3.3.2.1.1). Therefore for those living in flats or large houses the mains sensors might be placed outside the operating range for the information to be received by the electricity monitor.
3. For this data collection, the electricity consumption data was collected using a data logger with the recorded information downloaded, periodically, to a secure server. This means that the participants had to have a fixed Internet connection within their homes.

3.3.2.1. Whole house electricity consumption data collection

The collection of the whole house electricity consumption data involved the placing of three pieces of equipment into the participant's houses. These were the mains sensor, the electricity monitor and the data logger.

3.3.2.1.1. The mains sensor

The mains sensor has two parts, as shown in figure 3.1, the clip sensor and the transmitter. The clip sensor, as highlighted in figure 3.1, is clipped around the main electrical feed cable from the household's electricity meter to its fuse box. This sensor monitors the magnetic field generated around the mains cable to measure the current passing through it. The transmitter, as highlighted in figure 3.1, then transmits this information wirelessly to the electricity monitor (as described in section 3.3.2.1.2) at six-second time intervals.

Figure 3.1: An example of the mains sensor

3.3.2.1.2. The electricity monitor

The electricity monitor, as shown in figure 3.2, receives the information from the mains sensors (as described in section 3.3.2.1.1) and displays this as energy usage of the house, at that time, in Watts. The electricity monitor also displays other information about the current and previous electricity usage, for example the current cost (in pence) of the energy usage of the house as well as the total power usage over the past 24 hours, week and month. This monitor also records the current temperature of the room, where the electricity monitor is situated. This information is provided as a visual breakdown of the power usage of the house to inform the occupants of their electricity consumption usage; however, it is not relevant to this research.

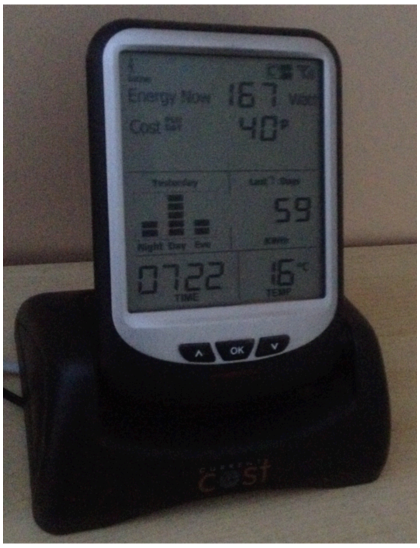

Figure 3.2: An example of the electricity monitor

3.3.2.1.3. The data logger

The data logger (as shown in figure 3.3) is plugged into the electricity monitor and is used to save the whole house electricity consumption data, as recorded and displayed by the electricity monitor. For this research the data logger is also connected to the Internet and the data are periodically downloaded to a secure web server and from there was accessed via a secure login. Section 3.3.2.3 gives more information about how the whole house electricity consumption data was access and subsequently manipulated to a form that can be used for further analysis.

Figure 3.3: An example of the data logger

3.3.2.2. Electricity diary data

The participants were asked to complete a record when they used any of the appliances listed below that they used for the week of the project.

The list of appliances that were recorded using the diaries were:

- Kettle
- Electric Oven
- Electric Hobs
- Television
- Washing machine
- Dishwasher
- Toaster
- Electric shower
- Microwave

The discussion provided in section 4.8, gives more detail into how the list of appliances were chosen based on the analysis of results from the survey (sections

4.3-4.7) and the review of the literature into the monitoring electricity consumption data (section 2.4).

As each household had different appliances and some used gas for cooking, the households were given some freedom in what they recorded. They were asked to record usage data (referred to here as the diary data) for the appliances on the list above (that they had in their house). For the households that used gas for cooking the occupants were also asked to record their usage of their extractor fan, with the aim of using this as a proxy for cooking in these households (as discussed in chapter 6).

Chapter 5 of this thesis will describe how the whole house electricity consumption data were combined with the usage diaries of the different electrical appliances to provide a set of training and test data for each, recorded, electrical appliance. These data were then used to train and test a model to recognise certain specific and/or general activities or tasks from a household's electricity consumption.

3.3.2.3. How the data was received and stored

The electricity consumption data were provided by the equipment as described in section 3.3.2.1 and then stored on a secure file server as a series of zipped text files. For the analysis of the electricity data, the zipped data files were unzipped and combined into one text file, using java-programming language.

The electricity consumption data were stored in the text file in the format shown below:

<time>1366711756</time><msg><src>CC128-v1.29</src><dsb>00365</dsb><time>10:52:38</time><tmpr>14.3</tmpr><sensor>0</sensor><id>03652</id><type>1</type><ch1><watts>00105</watts></ch1></msg>

Figure 3.4: Example of data recorded in text file

This data included the time in UNIX time, the identification number of the monitor, the temperature at that time (as recorded by the electricity monitor) and the energy consumption of the house at the time shown in Watts. Each line in the file represented a new-recorded data point.

To analyse the electricity data, certain elements of these data needed to be extracted. To do this a Matlab script file (as outlined in section 3.4.5) was created to read each line of the .txt file and place the relevant information into its own variable, for example, a separate variable for time, energy consumption and temperature. Once complete, the data could then be manipulated for further analysis. Section 5.2 of this thesis highlights in more details how these data were extracted and manipulated for the subsequent further analysis.

3.4. Electricity consumption data analysis

The previous section of the chapter (section 3.3.2) described the method used for this research to collected whole house electricity consumption data and appliance usage diary data. To achieve the aims of this research the collected data had to be analysed. Section 2.4.5.13 of the literature review provided an overview of the different methodologies and methods previously used by other researchers to analyse electricity consumption data and sensor data. This section will discuss the method chosen for how this electricity consumption data was analysed to achieve the aims of this research.

3.4.1. Feature set

Features (sometimes know as attributes, or variables in the social sciences) are the characteristics of an instance in the data (Witten et al., 2011).The choices of the features for use in this type of problem were very important because the features need to represent the different classes (in this study these are the electrical appliances), as well as finding features that were distinguishable between the different classes (appliances). The choice of the features were also important as poor choice in features will produce poor results for the model trained using the features (Theodoridis et al., 2010).

Features can take a number of different forms and four types can describe their structure, these are nominal, ordinal, interval and ratio. Nominal, or categorical, features, for example gender, are features that are distinguishable by the use of a name or label (Kantardzic, 2011). Nominal features have no numerical value and do not fit to any ordering scale (Witten et al., 2011). Ordinal features, for example the finishing positions in a race, are features that are possible to rank in some form of order (Kantardzic, 2011). Interval features, for example a measure of temperature in degrees Fahrenheit, are features that can be measure on a scale of fixed and equal units, although the origin is arbitrary (Kantardzic, 2011). Finally, ratio features, for example, age or weight of a person, are units of measurement where the origin is not arbitrary (Kantardzic, 2011). A feature set can incorporate features of different types, for example nominal and ordinal features (Witten et al., 2011).

3.4.1.1. Electricity consumption data- feature set

For this research the features were generated using the method reported by Lee et al., (2010) and Lin et al., (2010). Lee et al., (2010) and Lin et al., (2010) created a sliding window of the electricity data (as described below) to create features that were then used to train a dynamic Bayesian network to recognise when an electrical appliance had been turned on.

Lee et al., (2010) and Lin et al., (2010) created a sliding window of the data with a window size 7 (i.e., 7 sets of data taken at 5-second intervals, = 35 seconds) and then shifted the window by 1 sample. Once the sliding window had been created, it was then used to calculate features from the data. The features used in the work by Lee et al., (2010) and Lin et al., (2010) are shown below:

1. **Raw data**, i.e., the 7 data points for that time slice.
2. **Average**, i.e., the mean of the data points in that time slice.
3. **Peak value**, i.e., the maximum value of the data points in that time slice.
4. **Root mean square**, i.e., the root mean square of the data points in that time slice.
5. **Standard deviation**, i.e., the standard deviation of the data points in that time slice.
6. **Crest factor**, which for the paper by Lee et al., (2010) and Lin et al., (2010) is the window's peak value divided by the window's root mean square value.
7. **Form factor**, which for the paper by Lee et al., (2010) and Lin et al., (2010) is the window's root mean square value divided by the window's mean value.
8. **Peak to average ratio**, which for the paper by Lee et al., (2010) and Lin et al., (2010) is the window's peak value divided by the window's mean value.
9. **Delay ratio of the peak value,** which for the paper of Lee et al., (2010) and Lin et al., (2010) is calculated using the equation $\frac{1}{W} \times T_{Wh_{peak,t}}$ where W is the number of records in a sliding window and $T_{Wh_{peak,t}}$ is the index of the peak value within the window.

Section 5.4.2 of this thesis describes how the method of Lee et al., (2010) and Lin et al., (2010), which is described above was applied to the whole house electricity consumption data for this research.

3.4.2. Supervised and unsupervised learning

As discussed in more detail in section 2.4.5.5, there are two approaches to learning, i.e., supervised and unsupervised learning. For this research, as discussed in section 3.3.2, the electricity consumption data collected for this research consist of two parts, the whole house electricity consumption data and the appliance usage diary data. This provided a set of data (whole house electricity consumption data) and a corresponding set of targets (appliance usage diaries data), which provided a training set for supervised learning. Section 5.4.3 and 5.4.4 gives a description of how the electricity consumption data and the appliance usage data were combined and constructed to form the set of data and targets for the training set.

The supervised learning method chosen for this research was classification, because the aim was to identify when an appliance (from the list in section 3.3.2.2) had been used. The next section of this chapter will discuss the limitations of the data that affect the choice of the classification algorithm and the classification algorithm chosen for this research.

3.4.3. Classification method

As discussed in section 3.4.2, the supervised learning method chosen for this analysis was classification. The aim of this classification was to classify the whole house electricity consumption data into a number of classes (appliance usage), based on the training data provided. For classification the data can only be classified into the classes (appliances) that are present in the training data. This means that the model will only classify data into appliances for which it had been trained and that were present in the training set.

The structure of the whole house electricity consumption data produces a consideration with the choice of classifier algorithm for this research. The recording of electricity consumption data, at a frequency of every 6 seconds, produced a large amount of data points (91000+ for a week), although the usage of appliances, which were recorded by the usage diaries, for one house were 55 instances of appliance usage. This created an imbalance in the training set, as the majority of the data belonged to one class (the off class). Training with an imbalanced training set is usually described as a difficult task (Batista, Prati, & Monard, 2004) as generally classifiers require training with datasets, with an equal number of training points in each class.

Section 2.4.5 of the literature review highlighted different data mining and machine learning methods, which have previously been utilised by researchers for the classification of activities from sensor data. Support vector machines (SVM) used in the work of Fleury et al., (2010) were not considered for this analysis due to the imbalance of the data set. This issue was also raised by the work of Tang, Zhang, Chawla, & Krasser, (2009) who also noted that SVMs do not perform well on highly imbalanced datasets and produce a bias towards the majority class. Similarly, Artificial Neural Networks (Ruzzelli et al., 2010) were also not considered for this research due to their poor classification performance on highly imbalanced datasets (Mazurowski et al., 2008). Other examples such as Hidden Markov Models (Kröse et al., 2008; Singla et al., 2008) and decision rules (Farinaccio & Zmeureanu, 1999) are also sensitive to a class imbalance (Song, Morency, & Davis, 2013). The work of García, Fernández, & Herrera, (2009) and Liu, Chawla, Cieslak, & Chawla, (2010) highlighting that decision tree and rules are sensitive to class imbalance and can produce classifiers which are biased towards the majority class (Liu et al., 2010).

However, for this research, it is argued that the imbalance of the classes in the data is a reflection of the nature of electricity consumption data. As the recording of electricity consumption data represents people's habits and, as described by the work of Franco et al., (2008), with electricity usage people are generally habitual in their habits of appliance usage, although habitual to themselves. This should therefore be reflected in the choice of classifier.

For this research a probabilistic classifier that would take into account the residents habitual nature of appliance usage and their prior usage (prior probabilities) from the training data was chosen. The classifier for this research was a naïve Bayes classifier, which is described in more detail in the next section of the thesis.

3.4.4. Naïve Bayes classifier

For this research a naïve Bayes classifier was used to classify the whole house electricity consumption data into classes (appliance usage) based on the training data provided (as described in section 3.4.1). The implementation of the naïve Bayes classifier is shown in Chapter 5 along with the construction of the training and tests datasets, the feature sets and the interpretation of the results. This section will give an overview of the theory behind a naïve Bayes classifier and an example of how the naïve Bayes classifier is used to classify a sample feature set into a class, based on the training data.

A naïve Bayes classifier is probabilistic classifier based on Bayes' rule with some assumptions. The formula for Bayes' rule is shown below in equation 3.1.

$$P(Hypothesis|Data) = \frac{P(Data|Hypothesis)P(Hypothesis)}{P(Data)} \tag{3.1}$$

Where:

$P(Hypothesis|Data)$ is the proposed probability that the hypothesis is true given the observed data, this is call the posterior probability and is the result from the equation(Lee, 2004; Mitchell, 1997; Stone, 2013);

$P(Data|Hypothesis)$ is the probability of the data occurring based on the hypothesis (this is also called the likelihood (Lee, 2004; Mitchell, 1997; Stone, 2013);

$P(Hypothesis)$ is the probability of the hypothesis occurring based on prior knowledge, this is called the prior probability (Lee, 2004; Mitchell, 1997; Stone, 2013);

$P(Data)$ is the probability of the observed data, this is also called the marginal likelihood (Lee, 2004; Mitchell, 1997; Stone, 2013).

An easier way of showing the Bayes rule in shown below in equation 3.2 as described from the work of Stone, (2013):

$$Posterior = \frac{Likelihood \times Prior\ probability}{Marginal\ likelihood} \tag{3.2}$$

There are two assumptions that are assumed for the use of a naïve Bayes classifier; the first of these is that the features are independent of each other, given the class (i.e. that each feature contributes independently to the probability of a sample belonging to a particular class (Witten et al., 2011)). The second assumption is that for numerical features, the features within each class follow a normal distribution. However, in practice these assumption are either not possible to check or are violated, although as discussed by Soria, Garibaldi, Ambrogi, Biganzoli, & Ellis, (2011), Theodoridis & Koutroumbas, (2009) and Witten et al., (2011), even with the assumption violated, a naïve Bayes classifier still performs well.

To demonstrate how a naïve Bayes classifier is used to determine the posterior probability for this research, an example is given below. For this example, the naïve Bayes classifier will calculate the posterior probabilities of an appliance being turned on (i.e., belonging to the class microwave, washing machine, oven, dish washer or shower) or no appliance being turned on (i.e. belonging to the class off).

For this example the sample which is to be classified are numerical values, so a Gaussian probability density function (Theodoridis & Koutroumbas, 2009; Theodoridis et al., 2010) is used to calculate the probability of the sample data belonging to each class, based on the mean and standard deviation of that class, as calculated from the training data. The equation for a Gaussian probability density function is shown in equation 3.3.

$$P(x) = \frac{1}{\sqrt{2\pi\sigma^2}}\, exp\left(-\frac{(x-\ \mu)^2}{2\sigma^2}\right) \qquad (3.3)$$

where μ is the mean and σ is the standard deviation, calculated from the training data for each class.

For this study, there were six posterior probabilities to be calculated, these are the probabilities of the sample (table 3.1) belonging to one of the following classes, the dishwasher, the microwave, the washing machine, the oven, the shower and the off class (i.e. no recorded appliance was turned on). The sample was placed in the class that gave the highest posterior probability from equation 3.1. However, it is noted that the posterior probability of equation 3.1 is proportional to 1/P (Data). This means that as each equation is divided by the same marginal likelihood, the value of the posterior probability will change by the same proportion (Stone, 2013; Witten et al., 2011). This means that the marginal likelihood has no affect on the relative size of the posterior probabilities and therefore the class to which the sample is classified, so it is excluded from the equation.

The sample to be classified is shown in table 3.1.

	Average	Peak	Root mean square
Sample	2.583	9	3.905

Table 3.1: Example data to be classified

The prior probabilities of each of the classes are calculated from the training set, based on the number of instances of each of the classes (instances of appliance usage) in the training set divided by the total number of data points in the training set. The prior probabilities for each of the classes are shown in table 3.2.

	Off	Micro	Wash	Oven	Dish	Shower
Prior probability	0.999673491	9.60E-05	5.76E-05	3.84E-05	3.84E-05	9.60E-05

Table 3.2: Table of the prior probabilities of each class

To calculate the probability of the sample belonging to each of the classes, the Gaussian probability density function was used (equation 3.3). From equation 3.3 the mean and standard deviation, for each class had to be calculated. This was done by calculating the mean and standard deviation of each of the features, shown in the training data set for each of the six classes. The means and standard deviations for use in the Gaussian probability density function (equation 3.3) for each of the sample features, for each of the six classes are shown in table 3.3.

		Features		
		Average	**Peak**	**Root mean square**
Dishwasher	Mean	937.208	1917	1327.889
	Standard deviation	40.128	127.279	57.904
Microwave	Mean	601.500	1233	855.871
	Standard deviation	28.550	37.397	37.883
Oven	Mean	1024	2164.5	1496.708
	Standard deviation	24.042	166.170	77.606
Shower	Mean	4513.183	9145	6411.991
	Standard deviation	33.191	93.140	38.193
Washing machine	Mean	603.139	1503.333	945.020
	Standard deviation	93.696	262.169	111.999
Off	Mean	-0.602	60.132	60.429
	Standard deviation	193.657	385.090	299.475

Table 3.3: Means and standard deviations (for equation 3.3) for the classes based on the training data

The equations for calculating the posterior probabilities, *P*, for the sample (table 3.1) belonging to each of the classes are shown below.

$$P(Dishwasher) = P(Average|\ Dishwasher) \times P(Peak|\ Dishwasher) \times P(Root\ mean\ squares|\ Dishwasher) \times P(Dishwasher) \quad (3.4)$$

Where $P\ (Dishwasher)$ is the prior probability of the dishwasher as shown in table 3.2. To calculate $P(Average|\ Dishwasher)$ from equation 3.4, the Gaussian probability density function (equation 3.3) is used with the x in the equation being the value of the average given in the sample data (which, for this example, is 2.583). For this equation the μ is the mean and σ is the standard deviation of the training data (as shown in table 3.3) for the features. Thus, for this case, the mean and the standard deviation from the training data for the dishwasher (the average feature column which, for this case, is a mean of 937.208 and a standard deviation of 40.128). Putting these values into the Gaussian density function gives:

$$P(Average|\ Dishwasher) = \frac{1}{\sqrt{2 \times \pi \times (40.128)^2}} exp\left(-\frac{(2.583 - 937.208)^2}{2 \times (40.128)^2}\right)$$

$$= 1.5934\text{E-}120$$

This step is then repeated for $P(Peak|\ Dishwasher)$ and $P(Root\ mean\ square|\ Dishwasher)$. With the $P\ (Dishwasher)$ given from the prior probability (in table 3.2, for this case 3.84E-05) and the μ is the mean and σ is the standard deviation from the training data (table 3.3), therefore giving the final equation below.

$$P(Dishwasher) = 1.5934\mathrm{E}-120 \times 4.99833\mathrm{E}-52 \times 2\mathrm{E}-116 \times 3.84\mathrm{E}-05$$

$$= 6.26E-292$$

The remaining probabilities for each of the classes are calculated following the same steps with each of the probabilities for the remaining classes shown below.

$$P(Microwave) = 3.87084E-98 \times 2.5351\mathrm{E}-235 \times 1.5656\mathrm{E}-112 \times 9.60\mathrm{E}-05$$

$$= 0$$

$$P(Oven) = 0 \times 6.95614\mathrm{E}-40 \times 2.30836\mathrm{E}-83 \times 3.84\mathrm{E}-05$$

$$= 0$$

$$P(Shower) = 0 \times 0 \times 0 \times 9.60\mathrm{E}-05$$

$$= 0$$

$$P(Washing\ machine)$$
$$= 5.10629\mathrm{E}-12 \times 1.34114\mathrm{E}-10 \times 1.65662\mathrm{E}-18 \times 5.76\mathrm{E}-05$$

$$= 6.54E-44$$

$$P(Off) = 0.002059771 \times 0.001026879 \times 0.001308622 \times 0.999673491$$

$$= 2.76701\mathrm{E}-09$$

As the sample belongs to the class that gives the highest probability (i.e. highest value of all of the calculations), the sample data belongs to the off class.

Chapter 5 of this thesis describes how the electricity consumption data was pre-processed (section 5.2), transformed into a feature set (section 5.4.2) following the method of (Lee et al., 2010; Lin et al., 2010) as described in section 3.4.1 of this chapter, the extraction of the training and test datasets (section 5.4.3), the training of a naïve Bayes classifier (section 5.4.5) and test of this classifier with the analysis of the results from the test data set (section 5.4.6).

For this research, the analysis of the electricity consumption data was conducted in Matlab (as described in section 3.4.5) with naïve Bayes classifier created using an inbuilt Matlab function (as described in section 5.4.5).

3.4.5. Matlab

Matlab[2] is a suite of computer programs that provides a workspace for data analysis, data visualisation and programming. Matlab is a useful tool for analysing data, writing algorithms and model creation. Matlab also provides several in-built algorithms and toolboxes to help with the analysis of certain types of data for example signal processing or neural network design.[3]

The steps of how the electricity data were extracted, transformed into a feature set, spilt into training and test datasets and how the results from the naïve Bayes classifier are interpreted are shown in Chapter 5.

3.5. Conclusion

This chapter has provided an overview of the different methodologies and methods used to collect the three sets of data for this research. The next chapter in this thesis (Chapter 4) will show the analysis of the first set of collected data, the survey.

---

[2] http://uk.mathworks.com/products/matlab/features.html

[3] http://www.mathworks.co.uk/products/matlab/

Chapter 4: Survey Analysis and Results

4.1. Introduction

As shown in the literature review in Chapter 2, there is a body of research outlining the range of systems (for example, telecare), the range of activities to be measured and some research on the views of the people who are using monitoring technologies, to monitor certain aspects of their health or wellbeing. However, the literature review identified that only a limited amount of research has been undertaken on the key activities and features that carers/relatives, who are potential users of the information, would like to have access to, in order to be re-assured about their elderly/ill relative. Having identified this gap, a survey was developed with the aim of gaining further understanding of the views of relatives of an elderly or ill person into what types of activities should be monitored. The survey also assessed the views of relatives into the intrusiveness and required properties of any monitoring system, and who, in their view, should have access to the resulting information. The intention was that the results from this survey would inform which activities might be most useful to be monitored through measuring electricity consumption in the second stage of the study, described in chapters 5 and 6.

This chapter is divided into a number of different sections. Section 4.2 of this chapter provides an overview of the method used for the collection and analysis of the data from the survey. Section 4.3 presents the description of the responses from the survey. Section 4.4 presents the statistical analysis of the survey results. Section 4.5 provides a summary of the statistical survey results. Section 4.6 shows the results from the contents analysis conducted on three open-ended questions, which formed part of the survey. Section 4.7 presents the thematic analysis conducted on one of the open-ended questions that formed part of the survey. Finally, sections 4.8 and 4.9 present a discussion and a conclusion from the three-part analysis of the results from the survey.

4.2. Survey methods and data collection

For this part of the study, an online survey was created that contained a number of open and close-ended questions (as outlined in Chapter 3). To increase the response rate the survey was set up so that it could be completed by respondents who currently had an elderly or ill relative (current carers), by respondents who had previously had an elderly or ill relative (previous carers) or by respondents who had not had to look after an elderly or ill relative. The survey included a number of screening questions, which were designed so as to be able to differentiate between these three different groups of respondents and to tailor the questions to the respondent's previous experiences. These screening questions also formed part of the statistical analysis of the survey results as shown in section 4.4.

This survey was piloted on a small group of students at the University of Sheffield who were contacted via email. The email gave information about the survey and a link to the survey. The survey used for the pilot contained additional questions asking the respondents for their opinions about the survey. Improvements were made to the survey after reading the responses of this pilot.

The final online survey was distributed via an email to all those who were on a volunteer list at the University of Sheffield on the 19/10/2012. The email contained information about what the survey was about, how the data would be used and a link to the survey. For this survey no reminder email was sent out to those respondents who did not reply. This was due to the limitations of the volunteer mailing list used. Once the survey had been completed and the respondent had pressed the submit button the data were loaded into an online spreadsheet. The researcher via a secure login could access this spreadsheet and the data on it could be downloaded.

This survey received ethics approval from the Department of Computer Science Research Ethics Committee (as shown in appendix one).

4.2.1. Data analysis- statistics

For the statistical analysis of the responses the data were loaded into SPSS 20 and coded. Descriptive results of the survey responses are provided in section 4.3 with section 4.4 reporting the results of the statistical analysis of the survey responses.

4.2.2. Data analysis- content analysis

For the open-ended questions in the survey the textual responses were coded in SPSS 20 into different category headings based on what the respondent had written. From this, content analysis of all the responses was carried out and is shown in section 4.6.

4.2.3. Data analysis- thematic analysis

The written comments from the open-ended question in this survey were entered into NVivo 10. Using NVivo 10 the data were analysed to highlight themes in the responses given to the survey. The analysis of this is shown in section 4.7.

4.3. Survey results

The analyses of the responses from the survey were undertaken in a number of parts. This section will provide a description of the responses to the survey, in form of frequency tables based on the responses given to each of the questions of the survey.

Section 4.3.1 provides a description of the characteristics of the sample of the age and gender of those who responded to survey. Section 4.3.2 provides the characteristics of those for whom the respondents were currently caring. Section 4.3.3 provides the characteristics of those for whom the respondents had previously cared. Finally, section 4.3.4 provides the frequency of the responses to each of the main questions of this survey.

4.3.1. Characteristics of the sample

A total of 208 people responded to the survey, of whom 77.9% were female (n=162) and 22.1% were male (n=46). The age ranges of these participants are shown in table 4.1.

		Age Groups of Participants, n (%)					**Total**
		18-24	**25-30**	**31-40**	**41-50**	**50+**	
Gender of participants	**Male**	17 (23.3)	8 (25.0)	4 (13.8)	5 (20.0)	12 (24.5)	**46 (22.1)**
	Female	56 (76.7)	24(75.0)	25 (86.2)	20 (80.0)	37 (75.5)	**162 (77.9)**
Total		**73 (100)**	**32 (100)**	**29 (100)**	**25 (100)**	**49 (100)**	**208 (100)**

Table 4.1: Age and gender distribution of the participants of this survey

This survey asked whether the participant currently cared for an elderly or ill relative or had previously cared for an elderly or ill relative or had never cared for an elderly or ill relative. The numbers of participants in each group are shown in table 4.2. One half of the respondents had either cared for an elderly/ill relative in the past (30.3%) or were currently caring for someone (21.6%).

	Frequency (%)
Those who have never cared for an elderly or ill relative	100 (48.1)
Previously cared for an elderly or ill relative	63 (30.3)
Currently care for an elderly or ill relative	45 (21.6)
Total	**208 (100)**

Table 4.2: Distribution of participants within each caring group

For the respondents who were currently caring or had previously cared for an elderly or ill relative, questions were also asked about their relatives. The questions and the responses are shown in the sections 4.3.2 and 4.3.3 below respectively.

4.3.2. Respondents who currently cared for an elderly or ill relative

The distribution of the age and gender of the participant's relatives for whom they were currently providing care is shown in Table 4.3.

		Age group of relative being cared for, n							Total
		Not answered or missing	Under 50	50 -59	60-69	70-79	80-89	90+	
Gender of relative	**Male**	1	1	1	1	1	5	2	**12**
	Female	0	1	1	4	3	17	7	**33**
Total		**1**	**2**	**2**	**5**	**4**	**22**	**9**	**45**

Table 4.3: Distribution of the age ranges and gender of the participants' relatives (% values would not be meaningful here and are not included)

From Table 4.3 it can be seen that females comprised the larger gender group among the participants relatives, with n=33 (73.3%) of those who were currently looking after someone looking after a female relative. The age group with the most relatives was the 80-89 age group, in which n=17 (37.8%) were female and n=5 (11.1%) were male. Participants were asked whether their relative lived alone and whether the relative had a long-term illness or disease. Thirty-four participants (75.6%) reported that their relative had a long-term illness or disease and 29 (64.4%) reported that their relative lived alone. The approximate distance that the relative lived from the survey participant is shown in Table 4.4.

	Frequency (%)	Cumulative Percentage
Within 1 mile	11 (24.4)	24.40%
2-10 miles away	11 (24.2)	48.90%
11-100 miles away	13 (28.9)	77.80%
More than 100 miles but within the UK	5 (11.1)	88.90%
In a different country	5 (11.1)	100%
Total	**45 (100)**	

Table 4.4: Distribution of the distances from which the participants lived from their relatives

Table 4.4 shows the distances that the participants lived from their relative. It can be seen from Table 4.4 that almost half (48.9%) of the participants lived within 10 miles of their relative, i.e., within reasonable travelling distance by car.

4.3.3. Previously cared for an elderly or ill relative

For the 63 respondents who had previously cared for an elderly or ill relative, similar questions were asked about their relative. From these questions, 39 (61.9%) of the participants' relatives had been female, 48 (76.2%) had had a long-term illness or disease and 36 (57.1%) had lived alone.

	Frequency (%)	Cumulative Percentage
Within 1 mile	18 (28.6)	28.6%
2 -10 miles away	18 (28.6)	57.1%
11-100 miles	15 (23.8)	81%
More than 100 miles away but within the UK	9 (14.3)	95.2%
In a different country	3 (4.8)	100%
Total	**63 (100)**	

Table 4.5: Distribution of the distance from respondents that the relative lived

Table 4.5 shows the distance that the participants lived from their relative. From table 4.5 it can be seen that 57.1% of the participants (n=36) lived within 10 miles of their relative.

4.3.4. Main questionnaire statistics

In this section, the distributions of the responses to each of the questions asked in the main questionnaire are presented. The questions investigated a series of activities and asked respondents to rank and rate these activities based on whether they would want to be told that their relative has completed each of the activities. In this section, the overall results are reported, i.e., for the total sample. Table 4.6 shows the distribution of responses based on the respondents rating of each of the activities.

		Rating of each activity, n (row %)			
		Very Important	**Quite Important**	**Not at all Important**	**Total**
General Activities	**Changes in night time behaviour**	73 (35.1)	**112 (53.8)**	23 (11.1)	**208 (100)**
	Waking up	**100 (48.1)**	83 (39.9)	25 (12)	**208 (100)**
	Food Preparation	**122 (58.7)**	71 (34.1)	15 (7.2)	**208 (100)**
	Movement around the house	**114 (54.8)**	77 (37)	17 (8.2)	**208 (100)**
	Daytime general activities	85 (40.9)	**104 (50)**	19 (9.1)	**208 (100)**

Table 4.6: Distribution of the rating of each activity (**bold figures** indicate the modal response)

It can be seen from Table 4.6 that, of the activities, food preparation was most commonly described as very important 58.7% (n=122), changes in night-time behaviour was most commonly described as quite important with 53.8% (n=112) and waking up had the highest percentage of being not at all important, i.e., 12% (n=25).

Table 4.7 and 4.8 show the responses to the question asking the participants the most important and least important activity to be told that their relative had undertaken.

	Frequency (%)
Changes in night time behaviour	36 (17.3)
Waking up	46 (22.1)
Food preparation	**54 (26)**
Movement around the house	39 (18.8)
Daytime general activities	33 (15.9)
Total	**208 (100)**

Table 4.7: Distribution of the responses to the question "Please rank the most important activity to be told that your relative has completed"

From Table 4.7, it can be seen that the distribution of the most important activity varied from 15.9% (daytime general activities, n=33) to 26% (food preparation, n=54).

	Frequency (%)
Changes in night time behaviour	**65 (31.3)**
Waking up	50 (24)
Food preparation	20 (9.6)
Movement around the house	22 (10.6)
Daytime general activities	51 (24.5)
Total	**208 (100)**

Table 4.8: Distribution of responses to the question "Please rank the least important activity to be told that your relative has completed"

From Table 4.8, it can be seen that the distribution of the participants' views of the least important activity varied from 9.6% (food preparation, n=20) to 31.3% (changes in night time behaviour, n=65).

The next group of questions looked at the participants' views on whether they would find it important to know about specific activities their relative had performed. The distribution of responses to these questions is shown in Table 4.9.

		Responses to each activity, n (%)		
		No	**Yes**	**Total (%)**
Specific activities	**Used the Kettle**	74 (35.6)	**134 (64.4)**	**208 (100)**
	Watched TV	**122 (58.7)**	86 (41.3)	**208 (100)**
	Used the oven	69 (33.2)	**139 (66.8)**	**208 (100)**
	Used the washing machine	**118 (56.7)**	90 (43.3)	**208 (100)**
	Taken their medication	8 (3.8)	**200 (96.2)**	**208 (100)**

Table 4.9: Distribution of participants' responses to as to whether they wished to know whether their relative had undertaken specific activities

Participants were asked what type of activities (general or specific activities) they would want a remote monitoring system to record. The frequency of the responses is shown in Table 4.10.

	Frequency (%)
Both general and specific activities	**120 (57.7)**
General activities (e.g. that they are moving around the house)	57 (27.4)
Specific activities (e.g. that they turned the kettle on)	31 (14.9)
Total	**208 (100)**

Table 4.10: Distribution of the responses showing which type of activities participants want to be told that their relative has done

From Table 4.10, it can be seen that 120 participants (57.7%) wanted to be told that their relative had completed both general and specific activities.

The final question of the questionnaire asked the participant for their opinion on how important it was for a remote monitoring system to be non-intrusive. The distribution of the responses is in table 4.11.

	Frequency (%)
Very important	**137 (65.9)**
Quite important	58 (27.9)
Not at all important	13 (6.3)
Total	**208 (100)**

Table 4.11: Distribution of the responses to whether it is important for a remote monitoring system to be non-intrusive

From table 4.11, it can be seen that the majority of the participants (n=137; 65.9%) thought that it was very important for a remote monitoring system to be non-intrusive.

4.4. Statistical analysis of survey results

The next part of the analysis of the survey results involved the undertaking of Chi-squared tests (χ^2) to determine if there was any statistical association between the response of given to survey based on which age, gender or caring group the participant was in. For this analysis a significance level (α) of $p < 0.05$ was adopted.

4.4.1. Characteristics of the sample

As reported in section 4.3.1, of those who took part in the survey, 48.1% (n=100) had never cared for an elderly or ill relative before, 30.3% (n=63) had previously cared for an elderly or ill relative and 21.6% (n=45) currently cared for an elderly or ill relative. Table 4.12 shows the age groups of the participants in each caring group.

		Age groups of participants, n (%)					Total
		18-24	25-30	31-40	41-50	50+	
Caring groups	Those who have never cared for an elderly or ill relative	46 (63)	16 (50)	14 (48.3)	13 (52)	11 (22.4)	**100 (48.1)**
	Previously cared for an elderly or ill relative	19 (26)	10 (31.2)	9 (31)	8 (32)	17 (34.7)	**63 (30.3)**
	Currently care for an elderly or ill relative	8 (11)	6 (18.8)	6 (20.7)	4 (16)	21 (42.9)	**45 (21.6)**
Total		**73 (100)**	**32 (100)**	**29 (100)**	**25 (100)**	**49 (100)**	**208 (100)**

Table 4.12: Table showing the age ranges of participants in each caring group

There was a significant association between the age groups of participants and the caring group to which they belong (χ^2_{trend} = 15.86; df = 1; p < 0.001). It can be seen from table 4.12 that of the 73 participants in the age range 18-24 years old, 46 (63%) had never cared for an elderly or ill relative compared with 11 of the 49 people (22.4%) in the age group 50+. Conversely, of the 49 participants in the age group 50+, 21 (42.9%) were currently caring for a relative, compared with only eight of the 73 participants in the 18-24 years group (11%). There was a similar age-associated differential across participants who had previously cared for a relative. The gender breakdown of the different caring groups is shown in Table 4.13.

		Gender of Participants, n (%)		Total
		Male	Female	
Caring groups	Had never cared for an elderly or ill relative	20 (43.5)	80 (49.4)	**100 (48.1)**
	Previously cared for an elderly or ill relative	15 (32.6)	48 (29.6)	**63 (30.3)**
	Currently cared for an elderly or ill relative	11 (23.9)	34 (21.0)	**45 (21.6)**
Total		**46 (100)**	**162 (100)**	**208 (100)**

Table 4.13: Table showing the gender of participants in each caring group

There was not a significant association between the gender of participants and the caring group to which they belong (χ^2 = 0.506; df =2; p = 0.776). From Table 4.13 it can been seen that the highest percentage of both male (n=20; 43.5%) and female (n=80; 49.4%) belonged to the group who had never cared for an elderly or ill relative.

4.4.2. Main questions analyses

Table 4.14 shows the responses given and the chi-squared test results, based on the participants' care group, age and gender, in response to the question asking the participants to rate by importance, whether they would like to know that their relative had undertaken these activities.

		Gender, n (%)		Age group, n (%)					Caring group, n (%)		
		Male	**Female**	**18-24**	**25-30**	**31-40**	**41-50**	**50+**	**Never cared for an elderly or ill relative**	**Previously care for an elderly or ill relative**	**Currently care for an elderly or ill relative**
Changes in night time behaviour	**Very Important**	16 (34.8)	57 (35.2)	26 (35.6)	15 (46.9)	10 (34.5)	6 (24)	16 (32.7)	36 (36)	22 (34.9)	15 (33.3)
	Quite Important	**21 (45.7)**	**91 (56.2)**	**39 (53.4)**	**16 (50)**	**18 (62.1)**	**13 (52)**	**26 (53.1)**	**53 (53)**	**34 (54)**	**25 (55.5)**
	Not at all Important	9 (19.6)	14 (8.6)	8 (11)	1 (3.1)	1 (3.4)	6 (24)	7 (14.3)	11 (11)	7 (11.1)	5 (11.1)
	Total	46 (100)	162 (100)	73 (100)	32 (100)	29 (100)	25 (100)	49 (100)	100 (100)	63 (100)	45 (100)
	Chi-squared	χ^2= 4.604, df =2, p = 0.100		χ^2_{trend} = 1.639, df = 1, p = 0.201 *note: 3 cells (20%) have expected count less than 5*					χ^2= 0 .102, df = 4, p = 0.999		
Waking up	**Very Important**	16 (34.8)	**84 (51.9)**	30 (41.1)	**17 (53.1)**	14 (48.3)	**11 (44)**	**28 (57.1)**	43 (43)	**31 (49.2)**	**26 (57.8)**
	Quite Important	**21 (45.7)**	62(38.3)	**33 (45.2)**	13 (40.6)	**15 (51.7)**	**11 (44)**	11 (22.4)	**49 (49)**	23 (36.5)	11 (24.4)
	Not at all Important	9 (19.6)	16 (9.9)	10 (13.7)	2 (6.2)	0 (0)	3 (12)	10 (20.4)	8 (8)	9 (14.3)	8 (17.8)
	Total	46 (100)	162 (100)	73 (100)	32 (100)	29 (100)	25 (100)	49 (100)	100 (100)	63 (100)	45 (100)
	Chi-squared	χ^2 = 5.458, df = 2, p = 0.065		χ^2_{trend}= 0.311, df = 1, p = 0.577 *note: 3 cells (20%) have expected count less than 5*					χ^2= 9.239, df = 4, p = 0.055		
Food preparation	**Very Important**	20 (43.5)	**102 (63)**	**41 (56.2)**	**18 (56.2)**	**23 (79.3)**	**14 (56)**	**26 (53.1)**	**59 (59)**	**43 (68.3)**	**20 (44.4)**
	Quite Important	**21 (45.7)**	50 (39)	29 (39.7)	11 (34.4)	4 (13.8)	9 (36)	18 (36.7)	37 (37)	16 (25.4)	18 (40)
	Not at all Important	5 (10.9)	10 (6.2)	3 (4.1)	3 (9.4)	2 (6.9)	2 (8)	5 (10.2)	4 (4)	4 (6.3)	7(15.6)
	Total	46 (100)	162 (100)	73 (100)	32 (100)	29 (100)	25 (100)	49 (100)	100 (100)	63 (100)	45 (100)
	Chi-squared	χ^2= 5.710, df = 2, p = 0.058		χ^2_{trend}= 0.361, df = 1, p = *0.548* note*: 4 cells (26.7%) have expected count less than 5*					χ^2 = 10.484, df = 4, p = 0.033 note: 22.2% of cells have expected count less than 5		
Movement around the house	**Very Important**	**23 (50)**	**91 (56.2)**	**44 (60.3)**	**16 (50)**	**19 (65.5)**	12 (48)	**23 (46.9)**	**56 (56)**	**32 (50.8)**	**26 (57.8)**
	Quite Important	16 (34.8)	61 (37.7)	22 (30.1)	15 (46.9)	8 (27.6)	**13 (52)**	19 (38.8)	41 (41)	24 (38.1)	12 (26.7)
	Not at all Important	7 (15.2)	10 (6.2)	7 (9.6)	1 (3.1)	2 (6.9)	0(0)	7 (14.3)	3 (3)	7 (11.1)	7 (15.6)
	Total	46 (100)	162 (100)	73 (100)	32 (100)	29 (100)	25 (100)	49 (100)	100 (100)	63 (100)	45 (100)
	Chi-squared	χ^2 = 3.915, df = 2, p = 0.141		χ^2_{trend}= 1.704, df = 1, p = 0.192 *note: 4 cells (26.7%) have expected count less than 5*					χ^2= 8.975, d f= 4, p = 0.062		
Daytime general activities	**Very Important**	14 (30.4)	71 (43.8)	32 (43.8)	13 (40.6)	11 (37.9)	9 (36)	20 (40.8)	39 (39)	23(36.5)	**23 (51.1)**
	Quite Important	**23 (50)**	**81 (50)**	**34 (46.6)**	**18 (56.2)**	**17 (58.6)**	**14 (56)**	**21 (42.9)**	**55 (55)**	**32 (50.8)**	17 (37.8)
	Not at all Important	9 (19.6)	10 (6.2)	7 (9.6)	1 (3.1)	1 (3.4)	2 (8)	8 (16.3)	6 (6)	8 (12.7)	5 (11.1)
	Total	46 (100)	162 (100)	73 (100)	32 (100)	29 (100)	25 (100)	49 (100)	100 (100)	63 (100)	45 (100)
	Chi-squared	χ^2= 8.607, df =2, p = 0.014		χ^2_{trend}=0 .862, df = 1, p = 0.353 *note: 4 cells (26.7%) have expected count less than 5*					χ^2= 5.530, df = 4, p = 0.237		

Table 4.14: Distribution of participants caring groups, age and gender with respect to their response to rating of activities as well as chi-squared results for each.

From Table 4.14, it can be seen that there was a significant association between the gender of the participant and views on the importance of knowing about daytime general activities (*p* = 0.014). Although the importance of knowing about food preparation approached statistical significance (*p* = 0.058), it was not deemed significant. For the remainder of the activities there was no significant association between gender of participant and the importance of knowing about the activity. It can also be seen from the table that the importance of knowing about daytime general activities, movement around the house and changes in night-time behaviour was the same regardless of gender. For waking up and food preparation, there was a difference in results based on gender. The highest percentages among the female participants were for rating the activity most important for waking up (n = 84; 51.9 %) and for food preparation (n = 102; 63%). For the male participants the highest percentages were for rating the activity quite important, with n =21 (45%) both for waking up and for food preparation.

Analysing the data based on the age group of the participants, there were no significant associations between the age group of the participant and their view of the importance of knowing about each of the activities. However, there were some differences between the age groups and the responses given, but these may have arisen due to random variation. For example, for knowing about waking up, where for age groups 18-24 (n = 33; 45.2%) and 31-40 (n=15; 51.7%) the highest percentage rated the activity as quite important. For age groups 25-30 (n = 17; 53.1%) and 50+ (n=28; 57.1%) the highest percentage rated the activity very important. The age group 41-50 (n=11; 44%) had the same percentage for rating this as very and quite important. The other difference between age group and results was for knowing about movement around the house. The highest percentage for age group 41-50 (n=13; 52%) was for quite important but, for all the other age ranges, the highest percentage was for rating the activity as very important.

The final group on table 4.14 is caring group of the participants. There was a significant association between the type of caring group and views on the importance of knowing about food preparation (*p* = 0.033), although the expected cell count was less than 5 for 22.2% of these cells, which may have inflated the test statistic (Altman, 1999). A higher proportion of people who had previously cared for a relative (n=43; 68.3%) felt that knowing about food preparation was very important, compared to n = 59 (59%) people who had never cared for someone and n =20 (44%) for people who were currently caring for someone. The importance of knowing whether someone had woken up approached statistical significance (p = 0.055). Twenty-six of the 45 people who currently cared for a relative (57.8%) and 31 of the 63 (49.2%) people who had previously cared for someone felt that knowing about the person waking up was very important, compared with only 43 of the 100 (43%) people who had never cared for someone. For the remaining activities, there were no significant associations between caring group and the importance of knowing about the activity.

Table 4.15 and 4.16 show the responses given and chi-squared results based on the participants' gender, age and care group, in response to the question asking the participants the most important and least important activity to be told that their relative had undertaken.

		The most important activity to be told about, n (%)						
		Changes in night time behaviour	**Waking up**	**Food preparation**	**Movement around the house**	**Daytime general activities**	**Total**	**Chi-squared test**
Gender	**Male**	11 (23.9)	**12 (26.1)**	11(23.9)	5 (10.9)	7 (15.2)	46 (100)	χ^2**=3.977, df = 4, *p* = 0.409**
	Female	25 (15.4)	34 (21)	**43 (26.5)**	34 (21)	26 (16)	162 (100)	
Age	**18-24**	15 (20.5)	11 (15.1)	**18 (24.7)**	15 (20.5)	14 (19.2)	73 (100)	χ^2_{trend}**=0.168, df=1, *p* = 0.682[a]**
	25-30	7 (21.9)	**9 (28.1)**	6 (18.8)	6 (18.8)	4 (12.5)	32 (100)	
	31-40	5 (17.2)	3 (10.3)	**12 (41.4)**	6 (20.7)	3 (10.3)	29 (100)	
	41-50	2 (8)	**7 (28)**	6 (24)	5 (20)	5 (20)	25 (100)	
	50+	7 (14.3)	**16 (32.7)**	12 (24.5)	7 (14.3)	7 (14.3)	49 (100)	
Caring status	**Never cared for an elderly or ill relative**	20 (20)	17 (17)	**26 (26)**	19 (19)	18 (18)	100 (100)	χ^2**=9.103, df = 8, *p* = 0.334**
	Previously cared for an elderly or ill relative	10 (15.9)	14 (22.2)	**21 (33.3)**	10 (15.9)	8 (12.7)	63 (100)	
	Currently care for an elderly or ill relative	6 (13.3)	**15 (33.3)**	7 (15.6)	10 (22.2)	7 (15.6)	45 (100)	

Table 4.15: Distribution of participant's caring group in relation to responses to what is the most important activity to be told that your relative has done with chi-squared results for each ([a]16% of cells had expected count less than 5)

There was no significant association between gender, age or caring group and the most important activity to be aware that their relative had undertaken. From table 4.15, the most important activity to be aware that their relative had undertaken, for females, was food preparation (n=43; 26.5%) and for males waking up was the highest (n=12; 26.1%). For age groups 18-25 and 31-40, the most important activity was food preparation (n=18; 24.7%) and (n=12; 41.4%) respectively. For the other age ranges the most important activity was waking up for the 25-30 (n=9; 28.1%), 41-50 (n=7; 28%) and 50+ (n=16; 32.7%) age groups. The most important activity for those who had never cared and those who had previously cared was food preparation (n = 26; 26%) and (n = 21; 33%) respectively. For those who currently cared for someone, the most important activity to be told was waking up (n= 15; 33.3%).

		The least important activity to be told, n (%)						
		Changes in night time behaviour	**Waking up**	**Food preparation**	**Movement around the house**	**Daytime general activities**	**Total**	**Chi-squared test**
Gender	**Male**	12 (26.1)	**13 (28.3)**	3 (6.5)	6 (13)	12 (26.1)	46 (100)	χ^2=1.929, df=4, p = 0.749[a]
	Female	**53 (32.7)**	37 (22.8)	17 (10.5)	16 (9.9)	39 (24.1)	162 (100)	
Age	**18-24**	**18 (24.7)**	25 (34.2)	7 (9.6)	5 (6.8)	**18 (24.7)**	73 (100)	χ^2_{trend}= 0, df=1, p = 0.996[b]
	25-30	**11 (34.4)**	7 (21.9)	5 (15.6)	5 (15.6)	4 (12.5)	32 (100)	
	31-40	10 (34.5)	3 (10.3)	2 (6.9)	1 (3.4)	**13 (44.8)**	29 (100)	
	41-50	**10 (40)**	4 (16)	1 (4)	3 (12)	7 (28)	25 (100)	
	50+	**16 (32.7)**	11 (22.4)	5 (10.2)	8 (16.3)	9 (18.4)	49 (100)	
Caring status	**Never cared for an elderly or ill relative**	**36 (36)**	28 (28)	8 (8)	8 (8)	20 (20)	100 (100)	χ^2 =7.857, df=8, p = 0.448
	Previously cared for an elderly or ill relative	18 (28.6)	13 (20.6)	5 (7.9)	8 (12.7)	**19 (30.2)**	63 (100)	
	Currently care for an elderly or ill relative	11 (24.4)	9 (20)	7 (15.6)	6 (13.3)	**12 (26.7)**	45 (100)	

Table 4.16: Distribution of participants' caring group in relation to responses to what is the least important activity to be told that your relative has done with chi-squared results for each ([a] 20% of cells have expected count less than 5; [b]28% have expected count less than 5)

There was no significant association between gender, age or caring group and the least important activity to be aware that the participants' relative had done. From table 4.16, the least important activity for females was changes in night-time behaviour (n=53; 32.7%) and for males it was waking up (n=13; 28.3%). For age groups 25-30, 41-50 and 50+ the least important activity was changes in night-time behaviour with (n=11; 34.4%), (n=10; 40%) and (n=16; 32.7%) respectively. For age group 31-40 the least important was daytime general activities (n=13; 44.8%). The least important activities for the 18-24 year group were daytime general activities and changes in night-time behaviour (both n=18; 24.7%)). The least important activity for those who had never cared was changes in night-time behaviour (n=36; 36%). For those who had previously cared and currently cared, the least important activity was daytime general activities with (n=19; 30.2%) and (n=12; 26.7%) respectively.

Table 4.17 shows the responses given and the chi-squared test results based on the participant's gender, age and care group, in response to the question asking the participants to rate whether they would like to know that their relative had undertaken each of these activities.

		Gender, n (%)		Age, n (%)					Caring group, n (%)		
		Male	**Female**	**18-24**	**25-30**	**31-40**	**41-50**	**50+**	**Those who have never cared for an elderly or ill relative**	**Previously care for an elderly or ill relative**	**Currently care for an elderly or ill relative**
Used the kettle	**No**	18 (39.1)	56 (34.6)	27 (37)	**17 (53.1)**	9 (31)	7 (28)	14 (28.6)	32 (32)	23(36.5)	19 (42.2)
	Yes	**28 (60.9)**	**106 (65.4)**	**46 (63)**	15 (46.9)	**20 (69)**	**18 (72)**	**35 (71.4)**	**68 (68)**	**40 (63.5)**	**26 (57.8)**
	Total	46 (100)	162 (100)	73 (100)	32 (100)	29 (100)	25 (100)	49 (100)	100 (100)	63 (100)	45 (100)
	Chi-squared	χ^2 = 0.325, df = 1, p = 0.568		χ^2_{trend} = 2.210, df = 1, p = 0.137					χ^2 = 1.449, df = 2, p = 0.485		
Watched TV	**No**	**28 (60.9)**	**94 (58)**	**43 (58.9)**	**19 (59.4)**	**16 (55.2)**	**16 (64)**	**28 (57.1)**	**70 (70)**	31 (49.2)	21 (46.7)
	Yes	18 (39.1)	68 (42)	30 (41.1)	13 (40.6)	13 (44.8)	9 (36)	21 (42.9)	30 (30)	**32 (50.8)**	**24 (53.3)**
	Total	46 (100)	162 (100)	73 (100)	32 (100)	29 (100)	25 (100)	49 (100)	100 (100)	63 (100)	45 (100)
	Chi-squared	χ^2= 0.120, df = 1, p = 0.730		χ^2_{trend} = 0.004, df = 1, p = 0.948					χ^2 = 10.293, df = 2, p = 0.006		
Used the oven	**No**	17 (37)	52 (32.1)	17 (23.3)	11 (34.4)	7 (24.1)	10 (40)	24 (49)	24 (24)	23 (36.5)	22 (48.9)
	Yes	**29 (63)**	**110 (67.9)**	**56 (76.7)**	**21 (65.6)**	**22 (75.9)**	**15 (60)**	**25 (51)**	**76 (76)**	**40 (63.5)**	**23 (51.1)**
	Total	46 (100)	162 (100)	73 (100)	32 (100)	29 (100)	25 (100)	49 (100)	100 (100)	63 (100)	45 (100)
	Chi-squared	χ^2 = 0.381, df = 1, p = 0.537		χ^2_{trend} = 8.262, df = 1, p = 0.004					χ^2 = 9.125, df = 2, p = 0.010		
Used the washing machine	**No**	22 (47.8)	**96 (59.3)**	**37 (50.7)**	**16 (50)**	**19 (65.5)**	**16 (64)**	**30 (61.2)**	**55 (55)**	**34 (54)**	**29 (64.4)**
	Yes	**24 (52.2)**	66 (40.7)	36 (49.3)	**16 (50)**	10 (34.5)	9 (36)	19 (38.8)	45(45)	29 (46)	16 (35.6)
	Total	46 (100)	162 (100)	73 (100)	32 (100)	29 (100)	25 (100)	49 (100)	100 (100)	63 (100)	45 (100)
	Chi-squared	χ^2 = 1.908, df = 1, p = 0.167		χ^2_{trend} = 2.262, df = 1, p = 0.133					χ^2 = 1.409, df = 2, p = 0.494		
Taken their medication	**No**	0 (0)	8 (4.9)	1 (1.4)	0 (0)	2 (6.9)	1 (4)	4 (8.2)	1 (1)	1 (1.6)	6 (13.3)
	Yes	**46 (100)**	**154 (95.1)**	**72 (98.6)**	**32 (100)**	**27 (93.1)**	**24(96)**	**45 (91.8)**	**99 (99)**	**62 (98.4)**	**39 (86.7)**
	Total	46 (100)	162 (100)	73 (100)	32 (100)	29 (100)	25 (100)	49 (100)	100 (100)	63 (100)	45 (100)
	Chi-squared	χ^2=2.362, df = 1, p = 0.124[a]		χ^2_{trend} = 4.215, df = 1, p = 0.040[b]					χ^2 = 14.012, df = 2, p = 0.001[b]		

Table 4.17: Distribution of, chi-squared test results for, participant's caring group in relation to responses to knowing that their relative had done certain activities as well as each ([a] 25% of cells have expected count less than 5; [b]50% of cells have expected count less than 5)

From Table 4.17, it can be seen that there was no significant association between the gender of the participant and knowing whether their relative had performed certain activities. It can be seen from table 4.17 that, regardless of gender, the highest percentage of response given for each activity was the same except for knowing their relative had used the washing machine. For this activity, 96 of the 162 females (59.3%) did not want to know their relative had performed this activity. Conversely, 24 of the 46 males (52.2%) wanted to know that their relative had performed this activity.

Analysing the data in table 4.17, based on age group of the participants, there was a significant association between the age of participant and the importance of knowing whether their relative had used the oven (p = 0.004) and that they had taken their medication (p = 0.040). Note: for having taken their medication, 50% of cells had an expected count less than 5, which makes this finding less reliable, as low cell counts can artificially inflate the test statistic (Altman, 1999). For the remainder of the activities, there was no significant association between the age of the participant and the importance of knowing that their relative has completed certain activities.

From table 4.17, it can also be seen that, regardless of age of participant, the highest percentage of response given for each activity was the same except for knowing their relative had used the kettle. For this activity, 17 of the 32 people in the 25-30 year age group (53.1%) would not want to know that their relative had performed this activity. However, for all the other age groups they would want to know that their relative had performed this activity.

The final group on table 4.17 was the caring group of the participants. There was a significant association between the type of caring group and the importance of knowing whether their relative had watched TV (p = 0.006), used the oven (p = 0.010) and taken their medication (p = 0.001). Note: for having taken their medication, 50% of cells had an expected count less than 5 making this test less reliable. For the rest of the activities, there was no significant association between the caring group of the participant and the importance of knowing that their relative had undertaken certain activities.

The highest percentages of the responses given for each activity was the same regardless of the caring group of the participant except for knowing their relative had watched TV. For this activity 70 of the 100 respondents (70%) who had never cared for an elderly or ill relative did not want to know if their relative had watched TV, whereas a smaller proportion of those in the other two caring groups wanted to know that there relative had watched TV.

Table 4.18 shows the responses given and the chi-squared test statistic, based on the participants' care group, age and gender, in response to the question asking the participants what type of activities they want to know that their relative has done.

		Type of activities, n (%)				
		Both general and specific activities	**General activities (e.g. that they are moving around the house)**	**Specific activities (e.g. that they turned the kettle on)**	**Total**	**Chi-squared test**
Gender	**Male**	**22 (47.8)**	12 (26.1)	12 (26.1)	46 (100)	χ^2=5.990, df = 2, p = 0.05
	Female	**98 (60.5)**	45 (27.8)	19 (11.7)	162 (100)	
Age	**18-24**	**33 (45.2)**	28 (38.4)	12 (16.4)	73 (100)	χ^2_{trend} =1.062, df=1, p = 0.303[a]
	25-30	**22 (68.8)**	6 (18.8)	4 (12.5)	32 (100)	
	31-40	**20 (69)**	7 (24.1)	2 (6.9)	29 (100)	
	41-50	**15 (60)**	5 (20)	5 (20)	25 (100)	
	50+	**30 (61.2)**	11 (22.4)	8 (16.3)	49 (100)	
Caring status	**Those who have never cared for an elderly or ill relative**	**57 (57)**	28 (28)	15 (15)	100 (100)	χ^2= 0.664, df=4, p = 0.956
	Previously cared for an elderly or ill relative	**37 (58.7)**	18 (28.6)	8 (12.7)	63 (100)	
	Currently care for an elderly or ill relative	**26 (57.8)**	11 (24.4)	8 (17.)	45 (100)	

Table 4.18: Distribution of participant's age, gender and caring group according to responses to what types of activities do they would want to be told that their relative had done ([a] 3 cells (20%) have expected counts less than 5)

From table 4.18, it can be seen that there was a significant association between the gender of the participant and what types of activities the participants wanted to be told that their relative had done (p = 0.05). There was no significant association between age or caring group and the type of activities the participants wanted to be told that their relative had done.

Table 4.19 shows the distribution of responses and the results of the chi-squared tests, based on the participants care group, age and gender, in response to the question how important it was for a remote monitoring system to be non-intrusive.

<table>
<tr><th colspan="2" rowspan="2"></th><th colspan="5">Importance of a remote monitoring system being non-intrusive, n (%)</th></tr>
<tr><th>Very important</th><th>Quite important</th><th>Not at all important</th><th>Total</th><th>Chi-squared test</th></tr>
<tr><td rowspan="2">Gender</td><td>Male</td><td>29 (63)</td><td>14 (30.4)</td><td>3 (6.5)</td><td>46 (100)</td><td rowspan="2">χ^2=0.216, df=2, p = 0.898</td></tr>
<tr><td>Female</td><td>108 (66.7)</td><td>44 (27.2)</td><td>10 (6.2)</td><td>162 (100)</td></tr>
<tr><td rowspan="5">Age</td><td>18-24</td><td>48 (65.8)</td><td>22 (30.1)</td><td>3 (4.1)</td><td>73 (100)</td><td rowspan="5">χ^2_{trend} =0.139, df=1, p = 0.709 [a]</td></tr>
<tr><td>25-30</td><td>21 (65.6)</td><td>8 (25)</td><td>3 (9.4)</td><td>32 (100)</td></tr>
<tr><td>31-40</td><td>19 (65.5)</td><td>9 (31)</td><td>1 (3.4)</td><td>29 (100)</td></tr>
<tr><td>41-50</td><td>18 (72)</td><td>5 (20)</td><td>2 (8)</td><td>25 (100)</td></tr>
<tr><td>50+</td><td>31 (63.3)</td><td>14 (28.6)</td><td>4 (8.2)</td><td>49 (100)</td></tr>
<tr><td rowspan="3">Caring status</td><td>Those who have never cared for an elderly or ill relative</td><td>67 (67)</td><td>30(30)</td><td>3 (3)</td><td>100 (100)</td><td rowspan="3">χ^2= 4.119, df =4, p = 0.390 [b]</td></tr>
<tr><td>Previously cared for an elderly or ill relative</td><td>42 (66.7)</td><td>16 (25.4)</td><td>5 (7.9)</td><td>63 (100)</td></tr>
<tr><td>Currently care for an elderly or ill relative</td><td>28 (62.2)</td><td>12 (26.7)</td><td>5 (11.1)</td><td>45 (100)</td></tr>
</table>

Table 4.19: Table of participant's age, gender and caring group against responses to how important is it for a remote monitoring system to be non-intrusive with the chi-squared results for each ([a] 5 cells (33.3%) have expected count less than 5) ([b] 2 cells (22.2%) have expected count less than 5)

From table 4.19, it can be seen that there was no significant association between the participant's age, gender, caring status and the importance of a remote monitoring system being non-intrusive.

4.5. Summary of survey results

From the basic description of the survey results given in section 4.3.4, it can be seen that the majority of respondents class the general activities (table 4.6) as being either very important or quite important to be aware of. The highest percentage of respondents classed food preparation as the most important general activity to be aware of, changes in night-time behaviour was classed as the least important to be aware of. For the specific activities (as shown in table 4.9), the majority of the respondents wanted to know that their relative had undertaken three out of the five activities, these were using the kettle, using the oven and taking medication. Finally, the majority of the respondents wanted a remote monitoring system that showed both general and specific activities and was non-intrusive.

From the chi-squared tests presented in section 4.4.2, there were very few statistical associations between the caring group, age and gender of participants and the responses that they gave.

The responses for which there was a significant association between the caring group and knowing whether their relative had performed a certain activity were: whether their relative watched TV (p = 0.006) or whether their relative used the oven (p = 0.010). It is also noted that there were significant associations with knowing about taking medications (p = 0.001) and food preparation (p = 0.033), however the chi-squared results should be treated with caution for both of these results (as for medication 50% of cells have an expected count less than 5 and for food preparation the expected cell count was less than 5 for 22.2% of the cells).

There was a significant association between age group and knowing whether their relative had used the oven (p = 0.004). There was also a significant association knowing that their relative had taken their medication (p = 0.040) but the chi-squared result should be treated with caution as 50% of cells had an expected count less than 5.

There was significant association between gender and the importance of knowing about daytime general activities (p = 0.014) and what types of activities the participants want to be aware of that their relative has done (p = 0.05).

4.6. Content analysis

As part of the survey, three open-ended questions were asked. To analyse the responses given by the participants to these three questions contents analysis was used. Content analysis as described by Bryman, (2012) is a method of analysis of documents or text, so as to provide quantifiable content in terms of categories. For the responses given to the three open ended questions, similar responses were coded together into groups so as to provide frequencies for the responses.

Sections 4.6.1, 4.6.2 and 4.6.3 each provide the frequencies of the responses, based on the results of the contents analysis from each of the three open-ended questions. Each of these sections also provides a breakdown of the highest frequency responses for each of the caring groups (as highlighted in section 4.3.1).

4.6.1. Concerns raised by relatives

The first open-ended question asked in the survey was “Please list up to 3 events that most concern you, which may happen to your relative when they are alone”. The content analysis of all the responses given for this question is shown in table 4.20. The three most-frequently occurring responses from table 4.20 are shown in table 4.21.

		Frequency (%)
Concerns	Falls	154 (74)
	Not being able to call for help when needed	50 (24)
	Not eating or drinking properly	43 (20.7)
	Medical emergencies/ Sudden changes in health	39 (18.8)
	Accidents at home	25 (12.0)
	Injuring themselves	23 (11.1)
	Stroke	22 (10.6)
	Forgetting medication or not taking it correctly	21 (10.1)
	Feeling Lonely	20 (9.6)
	Heart Attack	20 (9.6)
	Bogus or unwanted callers (telephone or doorstep)	19 (9.1)
	Leaving appliances on	19 (9.1)
	Illness	18 (8.7)
	Being burgled	17 (8.2)
	Hurting themselves	17 (8.2)
	Not being able to get back up	13 (6.3)
	Becoming a victim of crime	11 (5.3)
	Fire	10 (4.8)
	Not being able to clean themselves	9 (4.3)
	Fainting	8 (3.8)
	Anxiety/ feeling scared	7 (3.4)
	Confusion	7 (3.4)
	Intruders	7 (3.4)
	Leaving the house and wandering off	7 (3.4)
	Struggling with day to day tasks	7 (3.4)
	Death	5 (2.4)
	Breathing problems	4 (1.9)
	Burns	4 (1.9)
	Depression	4 (1.9)
	Struggling to take care of themselves	4 (1.9)
	Assaults	3 (1.4)
	Being unable to do 'necessary things'	3 (1.4)
	Failure of house supply e.g. (heating)	3 (1.4)
	Feeling alone	3 (1.4)
	Incontinence	3 (1.4)
	Letting strangers into the house	3 (1.4)
	Being trapped in the house in the event of a fire	2 (1)
	Choking	2 (1)
	Feeling abandoned	2 (1)
	Feeling bored	2 (1)
	Low house temperature	2 (1)
	Not being able to move around the house	2 (1)
	Arguments with husband	1 (0.5)
	Being mistreated by their carer	1 (0.5)
	Being Mugged	1 (0.5)
	Being taken ill outside the home	1 (0.5)
	Being vulnerable	1 (0.5)
	Car accidents	1 (0.5)
	Cutting themselves	1 (0.5)
	Failure of their medical equipment	1 (0.5)
	Flooding	1 (0.5)
	Forgetting to lock doors	1 (0.5)
	Getting into hospital in time	1 (0.5)
	High blood pressure	1 (0.5)
	Negative comments from the public	1 (0.5)
	Not being able to get out of the bath	1 (0.5)
	Not being heard	1 (0.5)
	Not going to bed	1 (0.5)
	Not seeing people for weeks	1 (0.5)
	Psychological requirements	1 (0.5)
	Self-harming	1 (0.5)
	Their safety	1 (0.5)
	Unexpected letters or bills	1 (0.5)
	Using things that are unsafe for them	1 (0.5)

Table 4.20: Content analysis of responses given to concerns people have of things that could happen to their relative while they were alone

		Frequency (%)
Concerns	Falls	154 (74)
	Not being able to call for help when needed	50 (24)
	Not eating or drinking properly	43 (20.7)

Table 4.21: Table of the top 3 concerns of things that could happen to their relative while they are alone (n=208)

It can be seen from table 4.21 that the response with the highest percentage was falls, with n=154 out of the 208 respondents (74%) saying that this was one of their top three concerns. The second highest overall concern was that their relative would not be able to call for help when needed n=50 (24%). The third highest overall concern was that their relative was not eating or drinking properly (n =43; 20.7%).

4.6.1.1. Caring groups

The responses of the relatives' concerns were then separated into the three caring groups as shown earlier in section 4.3.1. The three most frequently occurring responses for each of these groups are shown in tables 4.22-4.24.

		Frequency (%)
Concerns	Falls	77(77)
	Not being able to call for help when needed	29 (29)
	Not eating or drinking properly	22 (22)

Table 4.22: Three most frequently occurring concerns of those who had never cared for an elderly or ill relative (n=100)

		Frequency (%)
Concerns	Falls	45 (71.4)
	Not being able to call for help when needed	15 (23.8)
	Not eating or drinking properly	15 (23.8)

Table 4.23: Three most frequently occurring concerns of those who had previously cared for an elderly or ill relative (n=63)

		Frequency (%)
Concerns	Falls	32 (71.1)
	Accidents at home	9 (20)
	Stroke	9 (20)

Table 4.24: Three most frequently occurring concerns of those who currently cared for an elderly or ill relative (n=45)

From the Tables 4.22-4.24 it can be seen that the response with the highest percentage across all three caring groups was falls with n =32 (71.1%) of those who currently cared for an elderly or ill relative, n=77 (77%) of those who had never cared for an elderly or ill relative and n=45 (71.4%) of those who had previously cared for an elderly or ill relative.

Similarly, it can also be seen that the second and third highest frequency of responses of those who had never cared for an elderly or ill relative (table 4.22) and those who had previously cared for an elderly or ill relative (table 4.23) were the same and correspond with the top 3 responses overall. However, in table 4.24, showing the top 3 responses of those who currently cared for an elderly or ill relative, the second and third highest frequency responses differed, with the joint second highest frequency being accidents at home and a stroke, both with n= 9 (20%) respondents.

4.6.2. Activities to give reassurance

The second open-ended question asked in the survey was "Please list up to 3 activities that would, if you knew that they had been undertaken, give you assurance of your relative's current status". The content analysis of all the responses given for this question is shown in table 4.25. The three most frequently occurring concerns from table 4.25 are shown in table 4.26.

		Frequency (%)
Activities to give reassurance	Regular visits	71 (34.1)
	Contact/Phone calls with relatives or carers	41 (19.7)
	Monitoring of food and drink consumption	32 (15.4)
	Someone to look after, care or assist them	25 (12)
	Emergency button or cord	22 (10.6)
	Food being prepared/ using of the cooker etc.	21 (10.1)
	Checks that medication has been taken or taken correctly	20 (9.6)
	Contact with neighbours or friends	20 (9.6)
	Monitoring of general household activities or activity levels	20 (9.6)
	Getting up	13 (6.3)
	Using the kettle	10 (4.8)
	Emergency fall alarm	7 (3.4)
	Knowing that they have left the house and returned safely	7 (3.4)
	Someone to keep them company/ live with them full time	7 (3.4)
	Using the toilet	7 (3.4)
	A message system giving current status	6 (2.9)
	Doors being locked or unlocked	6 (2.9)
	Going to bed/returning after getting up during the night	6 (2.9)
	Turning TV on/off	5 (2.4)
	Washing themselves/ Personal hygiene	5 (2.4)
	Webcam communications	5 (2.4)
	Friends or neighbours keeping an eye on them	4 (1.9)
	Getting dressed	4 (1.9)
	Monitoring to check that no appliances have been left on	4 (1.9)
	Drawing Curtains	3 (1.4)
	Food being delivered	3 (1.4)
	General health updates being sent to relatives or carers	3 (1.4)
	Being given things to do	3 (1.4)
	Reading	3 (1.4)
	Remote monitoring system	3 (1.4)
	A security system	3 (1.4)
	Someone to get to them quickly in an emergency	3 (1.4)
	Turning lights on/off	3 (1.4)
	Documents giving information on the relatives status	2 (1)
	Listening to the radio	2 (1)
	Picking up the post	2 (1)
	Reassurance that they are safe	2 (1)
	Reports on the homes security	2 (1)
	Stair lift	2 (1)
	Having more social contact with people	2 (1)
	Being more active	1 (0.5)
	Emergency alerts sent to the relative in an event of an emergency	1 (0.5)
	Filtering of phone calls	1 (0.5)
	Fire prevention	1 (0.5)
	Fitting locks to doors	1 (0.5)
	Gardening	1 (0.5)
	GPS tracking of the relative	1 (0.5)
	Health and safety assessment	1 (0.5)
	Alerts if the house temperature is too low	1 (0.5)
	I'm awake' button	1 (0.5)
	Knowing that they are wearing their personal alarm	1 (0.5)
	Doing light exercise	1 (0.5)
	Having a mobile phone that can be used by a partially sighted person	1 (0.5)
	Modifications to their house	1 (0.5)
	Seeing photos of what they have done during the day	1 (0.5)
	Stop smoking	1 (0.5)
	A telecom device	1 (0.5)
	A voice activating emergency alarm	1 (0.5)

Table 4.25: Content analysis of responses given to what activities would give you reassurance that your relative is well

		Frequency (%)
Activities	Regular visits	71 (34.1)
	Contact/Phone calls with relatives or carers	41 (19.7)
	Monitoring of food and drink consumption	32 (15.4)

Table 4.26: Top 3 responses given to what activities would give you reassurance that your relative is well (n=208)

From table 4.26 the activity with the highest frequency of responses was regular visits with (n=71; 34.1%). The second highest frequency was having contact or a phone call with the relative of the carers (n =41; 19.7%). The third highest frequency was monitoring of food and drink consumption (n =32; 15.4%).

4.6.2.1. Caring groups

The responses to what activities would give reassurance that their relative was well were separated according to the three caring groups as shown earlier in section 4.3.1. The three most frequently occurring responses for each of these groups are shown in tables 4.27-4.29.

		Frequency (%)
Activities	Regular visits	36 (36)
	Contact/Phone calls with relatives or carers	20 (20)
	Someone to look after, care or assist them	15 (15)

Table 4.27: Three most frequently occurring responses given to what activities would give them reassurance that their relative was well by those who had never cared for an elderly or ill relative (n=100)

		Frequency (%)
Activities	Regular visits	21 (33.3)
	Contact/Phone calls with relatives or carers	14 (22.2)
	Monitoring of food and drink consumption	10 (15.9)

Table 4.28: Three most frequently occurring responses given to what activities would give them reassurance that their relative was well by those who had previously cared for an elderly or ill relative (n=63)

		Frequency (%)
Activities	Regular visits	14 (31.3)
	Monitoring of food and drink consumption	11 (24.4)
	Contact/Phone calls with relatives or carers	7 (15.6)

Table 4.29: Three most frequently occurring responses given to what activities would give them reassurance that their relative was well by those who currently cared for an elderly or ill relative (n=45)

From tables 4.27-4.29 it can be seen that the responses with the highest frequency across all three caring groups was regular visits, with n =14 (31.3%) of those who currently cared for an elderly or ill relative, n=36 (36%) of those who had never cared for an elderly or ill relative and n=21 (33.3%) of those who had previously cared for an elderly or ill relative.

There were also other similarities between the tables with the second highest response for those who have never cared for an elderly or ill relative (table 4.27) and those who previously cared for an elderly or ill relative (table 4.28) was contact/phone calls with relatives or carers. This response was also identified by those who currently cared for an elderly or ill relative (table 4.29) but was the third most frequently occurring (n = 7; 15.6%).

For those who currently cared for an elderly or ill relative (table 4.29) the second highest response was monitoring of food and drink consumption with n =11 (24.4%). This response was the third most-frequently occurring for those who had previously cared for an elderly or ill relative with n =10 (15.9%).

From table 4.27, the third highest response for those who had never cared for an elderly or ill relative was someone to look after, care or assist them (n=15; 15%). This response was not among the highest three responses for the other two caring groups.

4.6.3. Properties of a remote monitoring system

The third open-ended question asked in the survey was “In your opinion, what properties does a remote monitoring system need to have”. The content analysis of all the responses given for this question is shown in table 4.30, with the three most-frequently occurring responses from table 4.30 shown in table 4.31.

		Frequency (%)
Properties	Timely reliable alerts or feedback to important parties	38 (18.3)
	Non-Intrusive/In obtrusive/Concealable	34 (16.3)
	Accurate and reliable	30 (14.4)
	Monitoring of activities around the house	25 (12)
	Easy, simple to use and understand	20 (9.6)
	Maintaining dignity and privacy	17 (8.2)
	Recording of sound and pictures	10 (4.8)
	Discreet	9 (4.3)
	Adaptable	7 (3.4)
	Doesn't disturb the person being monitored or force them to change their routine	7 (3.4)
	Easy to install/ not requiring any modifications to the house	7 (3.4)
	Emergency call button	6 (2.9)
	Ways of knowing that it is still working (long distance checks)	5 (2.4)
	Set up with the approval of those being monitored	5 (2.4)
	Economical	4 (1.9)
	Not to compensate for actual human contact	4 (1.9)
	Secure access	4 (1.9)
	Ability for the relative to disable the device if they don't want to be monitored	3 (1.4)
	Being able to tell who is performing the activities	3 (1.4)
	Fire/ carbon monoxide sensors	3 (1.4)
	Maintenance free	3 (1.4)
	Regular maintenance/ support	3 (1.4)
	Robust/ Can operate under all conditions	3 (1.4)
	To be intrusive	3 (1.4)
	Ability to access remotely to check on the relative	2 (1)
	Ability to be controlled remotely	2 (1)
	Accessible from a mobile device	2 (1)
	Anonymity	2 (1)
	Comprehensive	2 (1)
	Intelligent/Intuitive	2 (1)
	Monitoring of gas appliances- in case they are left on	2 (1)
	Reassuring to the relatives	2 (1)
	To be able to detect if the person needs help	2 (1)
	Battery Free	1 (0.5)
	Being able to talk to the relative	1 (0.5)
	Specific in its use	1 (0.5)
	Common sense	1 (0.5)
	Constant monitoring	1 (0.5)
	GPS/ movement tracking	1 (0.5)
	Having someone monitoring it	1 (0.5)
	Knowing that the activities have been done	1 (0.5)
	Making sure the relative knows that they aren't bothering anyone by using it	1 (0.5)
	Making the relative more active	1 (0.5)
	Provide daily medical information	1 (0.5)
	Provide peace of mind to the carers	1 (0.5)
	Security	1 (0.5)
	Small and compact	1 (0.5)
	Understanding of the residents shouts or screams	1 (0.5)
	Visual	1 (0.5)

Table 4.30: Content analysis of responses given to what properties does a remote monitoring system need to have

		Frequency (%)
Properties	Timely reliable alerts or feedback to important parties	38 (18.3)
	Non-Intrusive/unobtrusive/Concealable	34 (16.3)
	Accurate and reliable	30 (14.4)

Table 4.31: Top 3 responses given to what properties does a remote monitoring system need to have (n=208)

From table 4.31 the property with the highest frequency was timely reliable alerts or feedback to important parties with n=38 (18.3%) of respondents. The second highest frequency was non-Intrusive/unobtrusive/concealable with n=34 (16.3%) of respondents. The third highest frequency was it being accurate and reliable with n=30 (14.4%) of respondents. It can be seen that all three of the highest frequency responses had very similar frequencies/percentages.

4.6.3.1. Caring groups

The responses to the question on the properties that a remote monitoring system needed to have were separated into the three caring groups as shown earlier in section 4.3.1. The three most frequently occurring responses for each of these groups are shown in the tables below.

		Frequency (%)
Properties	Timely reliable alerts or feedback to important parties	21 (21)
	Non-Intrusive/In obtrusive/Concealable	17 (17)
	Accurate and reliable	14 (14)
	Monitoring of activities around the house	14 (14)

Table 4.32: Three most frequently occurring responses given to what properties does a remote monitoring system need to have by those who had never cared for an elderly or ill relative (n=100)

		Frequency (%)
Properties	Timely reliable alerts or feedback to important parties	12 (19)
	Monitoring of activities around the house	10 (15.9)
	Non-Intrusive/In obtrusive/Concealable	9 (14.3)

Table 4.33: Three most frequently occurring responses given to what properties does a remote monitoring system need to have by those who had previously cared for an elderly or ill relative (n=63)

		Frequency (%)
Properties	Accurate and reliable	9 (20)
	Non-Intrusive/In obtrusive/Concealable	8 (17.8)
	Easy, simple to use and understand	7 (15.6)

Table 4.34: Three most frequently occurring responses given to what properties does a remote monitoring system need to have by those who currently cared for an elderly or ill relative (n=45)

From tables 4.32 and 4.33, the highest frequency response for both those who had never cared and those who had previously cared for an elderly or ill relative was a timely reliable alerts or feedback to important parties, with n =21 (21%) from table 4.32 and n =12 (19%) from table 4.33. The highest frequency response for those who currently cared for an elderly or ill relative (table 4.34) was accurate and reliable with n =9 (20%). This was the same as the joint third highest response of those who had never cared for an elderly or ill relative (table 4.32).

The second highest response for both those who currently cared (table 4.34) and those who had never cared for an elderly or ill relative (table 4.32) was that a remote monitoring system should be non-Intrusive/unobtrusive/concealable with n=8 (17.8%) and n =17 (17%) responses respectively. Non-Intrusive/unobtrusive/concealable was the third highest response of those who had previously cared (table 4.33) with n =9 (14.3%).

From table 4.33, it can be seen that the second highest response of those who had previously cared for an elderly or ill relative was monitoring of activities around the house with n =10 (15.9%). This was also the join third highest response for those who had never cared for an elderly or ill relative (table 4.30) with n =14 (14%).

From table 4.34, the third highest response for those who currently cared for an elderly or ill relative was the remote monitoring being easy and simple to use and understand with n=7 (15.6%). This response was not in the three most frequently occurring responses for the other two caring groups.

4.7. Thematic analysis of relative’s concerns

As well as the contents analysis (as described in section 4.6), the first open-ended question (“Please list up to 3 events that most concern you, which may happen to your relative when they are alone”) asked to the participants was also analysed thematically. Thematic analysis as described by the work of Bryman, (2012) and Lapadat, (2010) is the extracting or identifying of themes from the data. To conduct the analysis the data was loaded into NVivo and the responses were coded into themes (Lapadat, 2010). This involved examining the responses given to this question and grouping the responses into recurring themes or sub themes that existed in the data. The themes and sub-themes, which were found from this thematic analysis, are discussed in the next section of this thesis.

4.7.1. Main theme and sub theme categories

From the conducting of the thematic analysis, into the response given by the respondents about their concerns for their elderly or ill relative, five main theme categories were found. These main themes were ‘ Accidents’, ‘Health’, ‘Security’, ‘Personal well-being’ and ‘Psychological health’. For each of these main themes there were also a number of subthemes, which are shown in figure 4.1.

The next five sections of this thesis (4.7.2- 4.7.6) discussed in more detail each of these main themes as well as their sub themes, with evidence provided from the responses to highlight and discuss each of these themes.

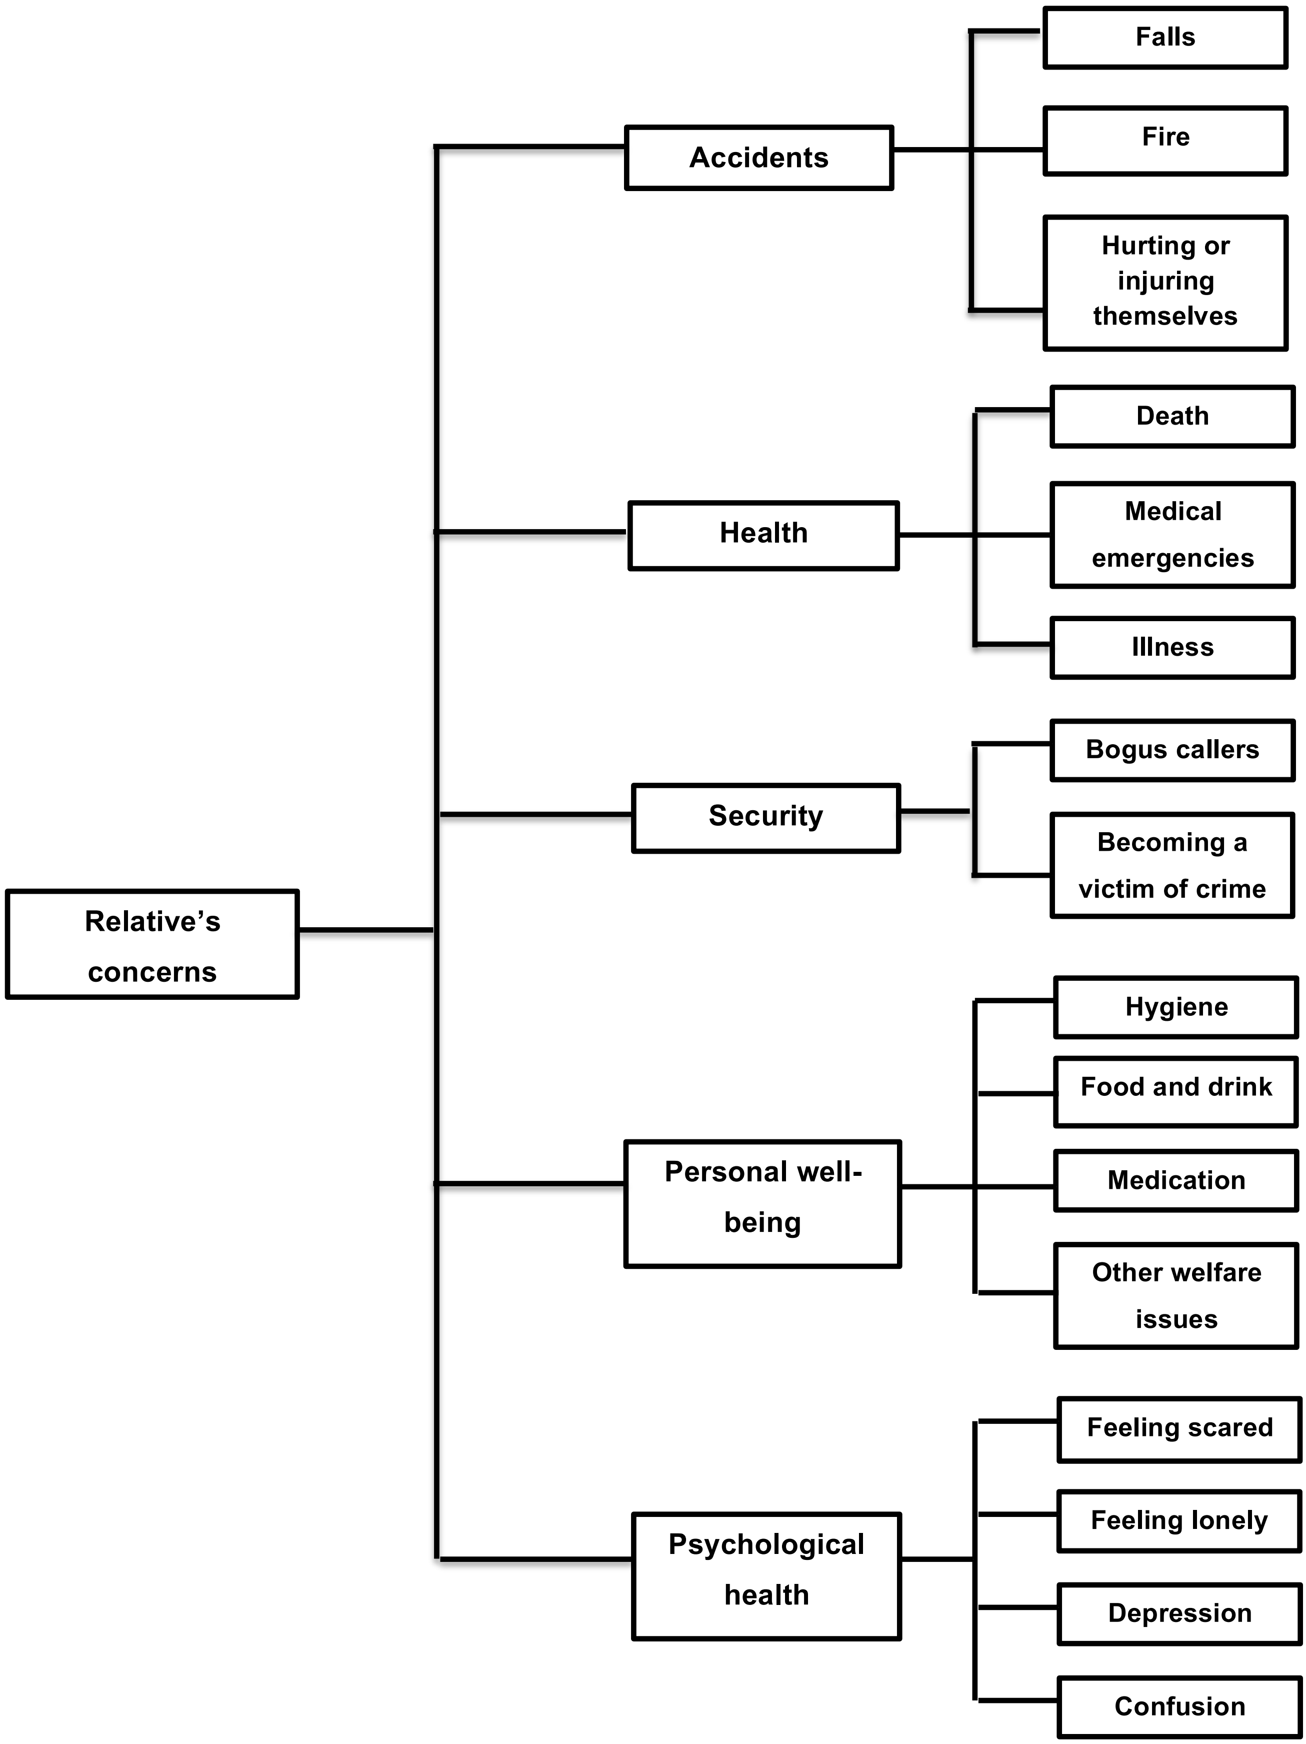

Figure 4.1: Figure showing the main themes categories and sub-theme categories of the relative's concerns

4.7.2. Accidents

The first main theme expressed by the respondents was their relative having an accident at home. The theme of accidents has been divided into a number of sub themes, relating to a range of different accidents or injuries.

4.7.2.1. Accidents- falls

Some participants expressed the concern about an elderly or ill relative falling:

Fall, falls, falling (mentioned by numerous respondents from all caring groups)

Falling down the stairs (Identification (ID) 9, Female, never cared for an elderly or ill relative)

If he wanted to go to the toilet he might fell down [sic] (ID 8, Female, Previously cared for an elderly or ill relative)

The concern was also expressed by some of the participants that it is not only the elderly or ill person falling that concerns them but that could also hurt or injure themselves because of the fall:

Fall and injure themselves (ID 5, Female, never cared for an elderly or ill relative)

When they are fell down and break their legs [sic] (ID 20, Female, Previously cared for an elderly or ill relative)

The concern that the elderly or ill relative would be left alone after a fall with no help is illustrated in the quotes below:

Falling over and not being able to get back up (mentioned by numerous respondents from all caring groups)

Fall & not have central call alarm to hand [sic] (ID 26, Female, Currently care for an elderly or ill relative)

Falling and being unable to contact me by telephone (ID 49, Female, Currently care for an elderly or ill relative)

Falling and no one knowing (ID 163, Female, never cared for an elderly or ill relative)

From these quotes it is not only the falling that concerns the respondents by also the consequences of the fall, both in terms of potential injury but also more importantly the inability to stand up or to raise an alarm.

4.7.2.2. Accidents- fire

Some participants expressed the concern of a fire within the elderly or ill relative's home:

A fire (ID 44, female, previously cared for an elderly or ill relative)

Setting fire to the house (ID 64, female, never cared for an elderly or ill relative)

The concern was also expressed that the elderly or ill relative could accidently cause the fire in their home:

Accidental fire from cooking (forgetting she has started cooking!) (ID42, female, currently care for an elderly or ill relative)

Cigarette not extinguished properly (ID26, female, currently care for an elderly or ill relative)

Danger with forgotten electrical/gas implements (ID59, male, never cared for an elderly or ill relative)

Leaving the gas on and causing a fire (ID107, female, never cared for an elderly or ill relative)

From these quotes it is not only the concern of a fire in the relative's house but also that they could accidentally cause the fire, for example while cooking, putting themselves in more danger.

4.7.2.3. Accidents- hurting or injuring themselves

The concern that the elderly or ill relatives would hurt or injure themselves is expressed in the quotes below:

That they might hurt themselves by accident (ID 28, female, never cared for an elderly or ill relative)

Personal injuries (ID17, male, currently cares for an elderly or ill relative)

Hurt themselves (mentioned by numerous respondents from all caring groups)

Injure or injuring themselves (mentioned by numerous respondents from all caring groups)

The concern of a more specific injuries or accidents happening to the elderly or ill relative is also illustrated in the quotes below:

Cooking burn / scald (ID29, Female, never cared for an elderly or ill relative)

They may choke on food/drink (ID 189, Female, Previously cared for an elderly or ill relative)

4.7.2.4. Accidents- summary

From the quotes highlighted in sections 4.7.2.1 to 4.7.2.3 it is clear that their relative having an unspecified accidents, injuries or hurting themselves while they are alone is a concern of the respondents. The concern was also expressed that performing certain everyday activities, such as cooking or eating, could be a cause for an accident or injury.

4.7.3. Health

The second main theme that was expressed by the respondents was concerns relating to the overall health or their relative. The theme of health was divided into a number of sub-themes relating to different health related concerns expressed by the respondents.

4.7.3.1. Health- death

Some participants expressed the concern of their elderly or ill relative dying:

Death (ID 193, male, currently care for an elderly or ill relative)

Some participants also expressed the concern of their elderly or ill relative dying alone:

Dying at home, alone (ID 206, female, currently care for an elderly or ill relative)

These quotes express not only the concern of the death or their relative but also that their relative is left to die alone.

4.7.3.2. Health- medical emergencies

Some participants expressed the concern that their elderly or ill relative could have a medical emergency or sudden change in their health:

Medical emergencies (ID 13, male, never cared for an elderly or ill relative)

Become suddenly unwell (ID 48, female, never cared for an elderly or ill relative)

Something requiring urgent medical assistance (ID194, female, never cared for an elderly or ill relative)

Unexpected sudden worsening of existing condition (ID23, female, currently care for an elderly or ill relative)

There was also the concern among participants that their elderly or ill relative could suffer from a specific medical emergency:

Stroke (mentioned by numerous respondents from all caring groups)

Heart attack (mentioned by numerous respondents from all caring groups)

Unable to breathe (ID114, female, previously cared for an elderly or ill relative)

Become unconscious (ID 61, male, never cared for an elderly or ill relative)

Fainting (ID 93, male, currently care for an elderly or ill relative)

The concern was also expressed that their elderly or ill relative would not be able to get help if they suffered a medical emergency:

Suffering a medical event (e.g. stroke, MI) and not being able to access care (ID 147, female, never cared for an elderly or ill relative)

That she gets so sick she is unable to telephone for help (ID179, female, previously cared for an elderly or ill relative)

Their condition worsens and they are unable to get help (ID157, female, previously cared for an elderly or ill relative)

In these quotes the respondents not only expressed their concerns about specific medical emergencies, for example, heart attacks, but also un-specific medical emergencies. There was also the concern among the respondents that if their relative was to suffer a medical emergency while they were alone they would not be able to get help.

4.7.3.3. Health- illness

Some participants expressed the concern that their elderly or ill relative might be suffering from an illness:

Illness (mentioned by numerous respondents from all caring groups)

Illness undiagnosed (ID 163, female, never cared for an elderly or ill relative)

There was also a concern that because of an illness they would be unable to take care of themselves properly:

Becoming ill and unable to get out of bed (ID49, female, currently cares for an elderly or ill relative)

Feeling too ill to look after herself (ie feed, wash etc) (ID 206, female, currently cares for an elderly or ill relative)

Feel too ill or unsteady to make a drink or get a meal (ID 56, female, previously cared for an elderly or ill relative)

4.7.3.4. Health- summary

From the concerns highlighted in sections 4.7.3.1 to 4.7.3.3, the respondents expressed concerns not only about a sudden change in the health of their relative but also the possibility of a slower change in health. The respondents identified that the change in health, either sudden or over time could cause their relative to be less able to care for themselves. There is also the concern that a sudden in change in health could leave their relative unable to call for help, leaving them alone.

4.7.4. Security

The third main theme that was highlighted by the respondents was a concern about the security of their relative. This main theme was divided into a number of sub-themes based on the different security concerns.

4.7.4.1. Security- becoming a victim of crime

The concern was expressed by some of the participants that their elderly or ill relative could become a victim of crime:

They may be the victim of criminals (ID 74, female, never cared for an elderly or ill relative)

Crime against them or their home (ID 85, female, never cared for an elderly or ill relative)

There was also a concern that they could a victim of a specific crime either against them or their property:

Burglary, being broken into, someone breaking in (mentioned by numerous respondents from all caring groups)

Intruder (ID 108, female, currently cares for an elderly or ill relative)

Assault (ID 105, male, previously cared for an elderly or ill relative)

Being attacked (ID 136, female, never cared for an elderly or ill relative)

That she gets mugged (ID 58, female, currently cares for an elderly or ill relative)

There was also a concern expressed by some of the participants at the repetition of crimes against their relative or their property:

Burgled again (ID 26, female, currently care for an elderly or ill relative)

Robbery – this has happen in last year (ID 37, female, never cared for an elderly or ill relative)

The concerns was also expressed by some participants that their elderly or ill relative could be seen as an ‘easy target’ due to their age or frailty:

That he would be attacked in the street as an easy target (vulnerable old man) (ID 129, female, never cared for an elderly or ill relative)

Someone might come and take advantage of their state (ID80, female, never cared for an elderly or ill relative)

These quotes not only express the concern that the participants had about their relatives’ becoming a victim of a specific or unspecific crime against them and their home, but also the concern that as they are elderly or ill that could be seen as an 'easy target' for criminals to be taken advantage of.

4.7.4.2. Security- bogus callers

There was a concern expressed by some of the participants about the kinds of people that call at their relative's door or call them on the telephone:

Unwanted visitors/telephone calls (ID31, female, previously cared for an elderly or ill relative)

Undesirable callers at the door (ID71, female, previously cared for an elderly or ill relative)

Someone calling at the door (ID177, female, currently cares for an elderly or ill relative)

People visiting the house that are trying to defraud my relative (ID 183, female, never cared for an elderly or ill relative)

Allowing a stranger into the house (ID 11, female, previously cared for an elderly or ill relative)

There was also the concern expressed by the participants that their relative's age or condition made them more vulnerable to these sorts of people:

He is vulnerable to malicious door-to-door sales people who could easily persuade him to sign up for a scam (ID 53, female currently care for an elderly or ill relative)

Preyed on by outside individuals, eg cold callers (ID 54, female currently care for an elderly or ill relative)

4.7.4.3. Security- summary

The quotes highlighted in section 4.7.4.1 and 4.7.4.2 have expressed the concerns of the respondents that their relatives are seen as more vulnerable and could be seen as a target for people to attack both inside and outside the home. There was also the concern that their relatives could fall victim to criminals from within their home via bogus visits or phone calls.

4.7.5. Personal well-being

The fourth main theme that was expressed by the respondents was concerns about the overall well-being of their relative. This main theme has been divided into a number of different sub-themes highlighting the many different concerns relating to the personal well-being of their relative.

4.7.5.1. Personal well-being- hygiene

Some of the participants expressed the concern that their elderly or ill relative could have problems with hygiene:

Being left in an insalubrious state as a result of not being able to physically attend to themselves (ID 78, male, never cared for an elderly or ill relative)

Personal hygine problems [sic] (ID 202, female, never cared for an elderly or ill relative)

Another concern expressed by some of the participants was the concern that their relative would not be able to complete personal or hygiene tasks if assistance was not there for them:

Need help in the shower (ID 140, female, currently cares for an elderly or ill relative)

Not being able to get to the toilet in time without help (ID 200, female, never cared for an elderly or ill relative)

There was also a concern among some of the participants about problems that could cause a hygiene issue with their relatives:

Not reaching bathroom in time to go to toilet (ID 186, male, currently cares for an elderly or ill relative)

Episodes of faecal incontinence (ID 71, female, previously cared for an elderly or ill relative)

These quotes express the concern among the participants that their relatives could suffer from problems with personal hygiene or problems that could lead to a person hygiene issue (for example incontinence). There was also the concern that their relative would struggle to maintain their standard of personal hygiene if their assistance was not there.

4.7.5.2. Personal well-being- food and drink

The concern was expressed by some of the participants that their elderly or ill relative would struggle with their food and drink intake:

Unable to get themselves food or drink (ID 180, female, never cared for an elderly or ill relative)

Difficulty with Food/Water consumption (ID13, male, never cared for an elderly or ill relative)

The concern was also expressed by some of the participants that their relative would forget to eat or prepare meals:

That they might not remember to eat (ID152, male, never cared for an elderly or ill relative)

Forgetting to eat and make meals (ID205, female, previously cared for an elderly or ill relative)

Another concern expressed by some of the participants was that their relative would not drink or eat properly or not eat or drink at all:

Not eating (ID 12, female, never cared for an elderly or ill relative)

Not eating or drinking enough (ID173, female, previously cared for an elderly or ill relative)

Not drinking properly (ID186, male, currently care for an elderly or ill relative)

Eating properly (ID184, female, previously cared for an elderly or ill relative)

The concern was also expressed by some of the participants that their relative could eat food that was bad for them:

Eating enough or food that is going off (ID172, female, currently cares for an elderly or ill relative)

Eat toxic food (ID 86, female, never cared for an elderly or ill relative)

The quotes above express the many concerns raised by the participants that their relative would not eat or drink sufficiently. This might be that due to the relative being unable to get food or drink for themselves, forgetting to eat or drink, just not eating or drinking or eating food that had gone off.

4.7.5.3. Personal well-being- medication

There was the concern expressed by some of the participants that their relative would forget to take their medication:

Failing to take medication (ID 160. female, currently cares for an elderly or ill relative)

Forgetting their medication (ID 185, female, never cared for an elderly or ill relative)

The concern was also expressed by some of the participant that their relative would not take their medication at the correct time:

Administering of medicines at right time (ID 137, male, previously cared for an elderly or ill relative)

They may forget to eat the pills on time (ID 20, female, previously cared for an elderly or ill relative)

There was also a concern among some of the participants that their relative would take too much of their medication:

Accidental overdose of prescribed medication (ID 173, female, previously cared for an elderly or ill relative)

Taking too much medication (ID 195, female, currently cares for an elderly or ill relative)

These quotes expresses the concern that the participants have that their relatives could forget to take their medication, they could take their medication at the wrong times of the day or they could take too much medication and overdose.

4.7.5.4. Personal well-being- other welfare issues

The concern was expressed by some of the participants about their elderly or ill relative leaving their house and wandering off:

They wander off outside alone (ID 22, female, previously cared for an elderly or ill relative)

Wander away from home and be unable to gind way back [sic] (ID 123, male, previously cared for an elderly or ill relative)

There was also a concern expressed by some of the participants that their relative could struggle with looking after themselves:

Struggling with everyday tasks (ID 112, female, previously cared for an elderly or ill relative)

They might struggle with day-to-day life (ID 121, male never cared for an elderly or ill relative)

Taking care of themselves properly (ID 191, female, never cared for an elderly or ill relative)

These quotes highlight some of the other welfare concerns that were expressed by the respondents.

4.7.5.5. Personal well-being- summary

The concerns highlighted in sections 4.7.5.1 to 4.7.5.4, have shown that the respondents expressed a range of concerns relating to the personal well-being of their relative. Amongst these are concerns about the ability to perform certain daily task such as hygiene and eating and drinking.

4.7.6. Psychological health

The final main theme expressed by the respondents was a concern for the psychological health of their relative. This theme was divided into a number of smaller sub-themes based on the specific concerns expressed by the respondents.

4.7.6.1. Psychological health- feeling scared

There was a concern expressed by some of the respondents that their elderly or ill relative would feel frightened or scared:

That they become frightened (ID 55, female, never cared for an elderly or ill relative)

Them being scared (ID 144, female, never cared for an elderly or ill relative)

These quotes expressed the concern raised by some of the respondents about how their relative feels when they are alone and the concern of them being scared or frightened.

4.7.6.2. Psychological health- feeling lonely

The concern expressed by some of the respondents that their relative could feel lonely:

Loneliness or feeling lonely (mentioned by numerous respondents from all caring groups)

A concern was also expressed by the respondents that they elderly or ill relative could feel alone or abandoned:

They feel alone/abandoned (ID 148, male, never cared for an elderly or ill relative)

Have a feeling of being abandoned (ID 155, female, previously cared for an elderly or ill relative)

Her feeling isolated and alone (ID 206, female, currently cares for an elderly or ill relative)

These quotes express the concerns of the respondents about how their relatives feel when they are alone and the concern that they could be left feeling not only lonely but also abandoned or isolated.

4.7.6.3. Psychological health- depression

A concern was expressed by some of the participants about their elderly or ill relative getting depressed:

Depression (mentioned by numerous respondents from all caring groups)

This quote expresses the concern of some of the respondents of their elderly or ill relative suffering from depression.

4.7.6.4. Psychological health- confusion

A concern was expressed by some of the participants about their elderly or ill relative become confused:

Become confused, confusion (mentioned by numerous respondents from all caring groups)

This quote expresses the concern of some of the respondents of their elderly or ill relative suffering from confusion.

4.7.6.5. Psychological health- summary

The sub-themes highlighted in section 4.7.6.1 to 4.7.6.4 have shown the concerns expressed by the respondents of how being alone, can lead to concerns that their relatives could suffer from a number of psychological health conditions.

4.7.7. Summary

This section has highlighted the themes and sub-themes in the responses collected from the first question of the survey. From this thematic analysis it is clear that the concerns of the respondents for their relatives cover many different aspects of their lives. As well as highlighting concerns of a range of different accidents or illnesses that could happen to their relative, there were also concerns that the relative could unintentionally cause himself or herself harm. There was also a concern highlighted in many of the themes that if something were to happen to their relative they would be unable to call for help or to continue to be able to look after themselves. The results from this section will be drawn together along with the results from the statistical analysis (section 4.4.) and the content analysis (section 4.6) in the next section to form a discussion of the results from the survey.

4.8. Discussion

4.8.1. Introduction

This section provides a discussion of the result from the three-stage analysis of the survey, as shown in sections 4.3 to 4.7. This discussion outlines how the results from this survey were subsequently incorporated into the range of appliances that were used for the collection of the electricity consumption data undertaken in chapters 5 and 6.

4.8.2. Monitoring of activities

From the analysis of the quantitative survey results (sections 4.3 and 4.4) as well as the content analysis (section 4.6), it is clearly important for the respondents that a remote system needs to be able to provide a relative or carer with information about food and drink consumption. This is demonstrated by the results shown in section 4.3.4, in which food preparation was the highest ranked important activity rated by the respondents in the survey. For more specific activities clearly linked with food and drink preparation, 64.4% of the respondents wanted to know that their relative had used the kettle, with 66.8% wanting to know that their relative had used the oven. Food and drink consumption also featured in the concern of the respondents, with 20.7% naming not eating or drinking as a concern to them (as highlighted in section 4.6.1).

Although food and drinking appeared widely across the responses in the survey, it was clear that a range of activities is also classed as important. More than 80% of the respondents to the survey classed all of the general activities in the survey as either very important or quite important, as shown in figure 4.2.

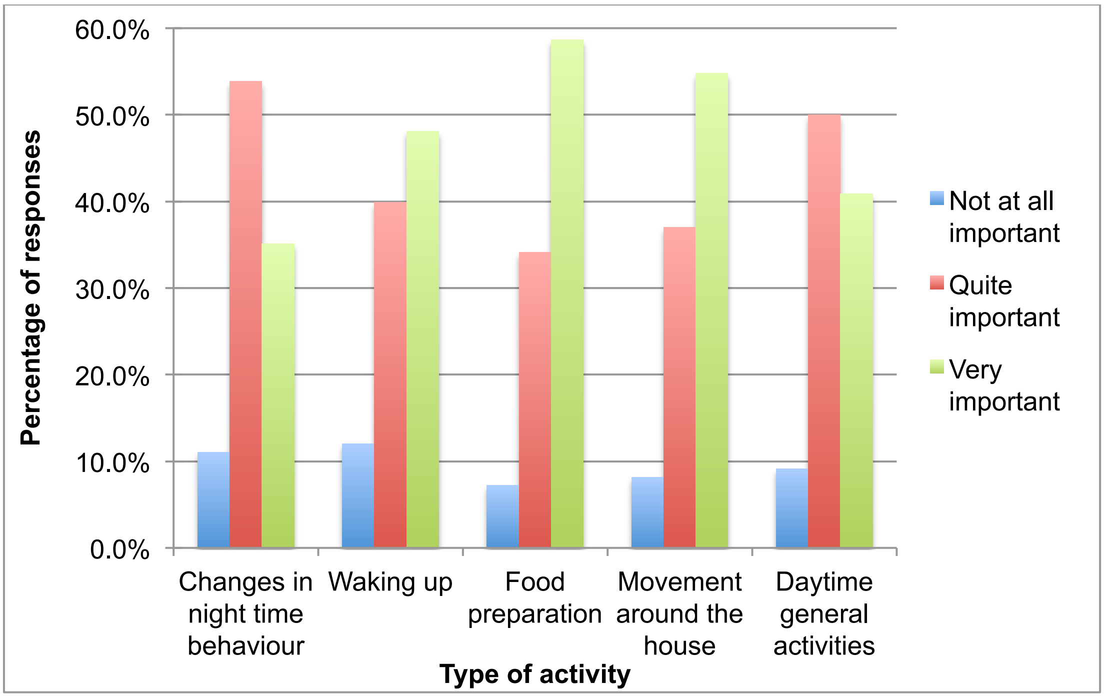

Figure 4.2: Frequency of responses indicating the importance of knowing certain general activities

The importance of showing a series of activities rather than focusing on just one activity has also been highlighted by the results presented in figure 4.3. Food preparation was classed by the highest percentage of respondents as the most important activity, with 26%. The second highest most important activity was waking up, with 22% of the respondents classing this as the most important. From figure 4.3, there is a difference of only 10% between the activity ranked as the most important by the highest percentage (26%) of respondents and the activity ranked the most important by the lowest percentage (16%) of respondents. This suggests that there is not a great variability in the level of importance across the sample.

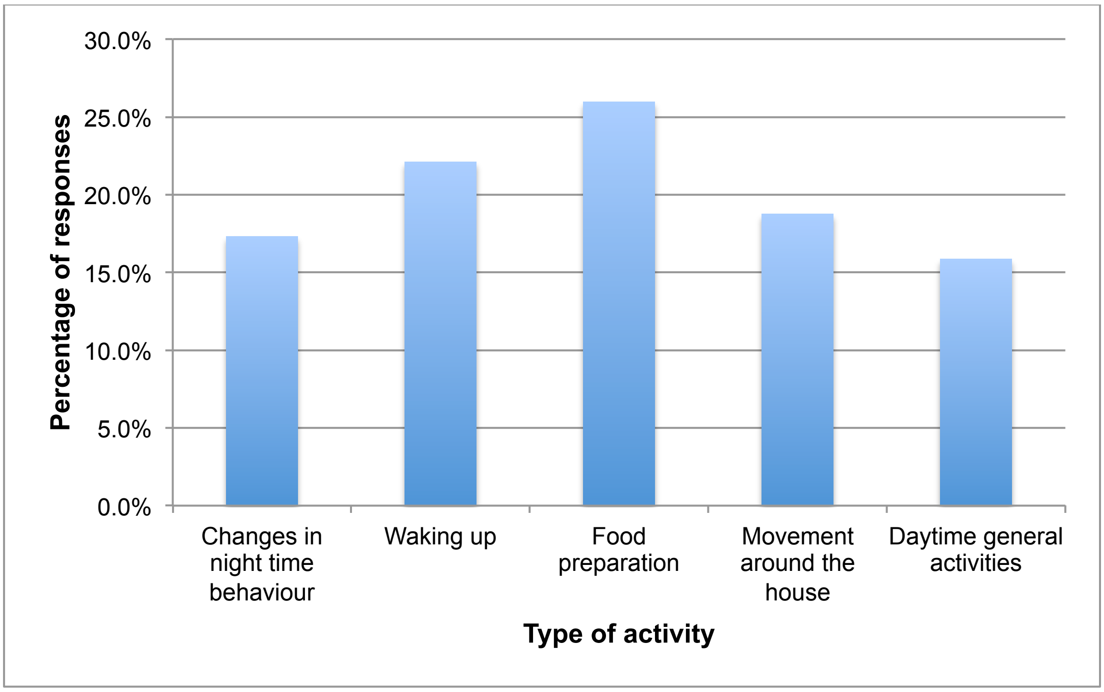

Figure 4.3: Figure showing the responses to the most important activity to be told from the survey

The results highlighted in figures 4.2 and 4.3 show that carers and relatives want to be told that a range of activities rather than a single activity or action had occurred. The monitoring of a series of activities may better reflect an assurance of normal daily living in the views of relatives and carers and changes over a number of activities or task may better indicate a longer term decline in health and well being. The views of the respondents that a remote monitoring system should monitor a range of activities is further supported by 57.7% of the respondents wanting to know a range of different activities rather than just a series of specific activities (as shown in table 4.9) or general activities (as shown in table 4.6).

4.8.3. Privacy and intrusiveness

Section 2.3.8 of the literature review highlighted the current research surrounding the privacy and ethical use of sensor technologies within homes. From this, the work of Bowes et al., (2012) and Stowe & Harding, (2010) highlighted that perceived intrusiveness and placing of sensors within homes can have an affect on the uptake of the use of monitoring systems. From the results to this survey, the need for a remote monitoring system to be non-intrusive was also important to the respondents, with 93.8% of the respondents classing the need for a remote monitoring system to be non-intrusive as either a 'very important' or 'quite important' property of a remote monitoring system. The need for a remote monitoring system to be non-intrusive was also highlighted by the content analysis, undertaken in section 4.6.3. From the content analysis, 16.3% of the respondents mentioned non-intrusiveness as a property of a remote monitoring system. As well as being non-intrusive, the respondents also highlighted other properties of a remote monitoring system as providing reliable and accurate feedback (18.3%) and overall reliability and accuracy (14.4%).

The privacy of data collected from households and who should have access to the data was also highlighted in the results from the survey, with 92.8% of the respondents indicating that carers and relatives should have access to the data produced from a remote monitoring system. From the responses 35.6% of the respondents also indicated that all relevant health care professionals should have access to the data with 13.9% of the respondents indicating that social services should have access to the data.

From the literature review in section 2.3.8, the work of Chan et al., (2008), Draper & Sorell, (2013), Milligan et al., (2011) and Stowe & Harding, (2010) have highlighted the fear that the use of remote monitoring technologies could lead to a reduction in the contact with their relatives or care providers. The results from this survey highlighted that, in the views of the respondents, the activity that would give them reassurance that their relative was well, had regular visits (34.1%) and regular contact with their relative or carers (19.7%). These responses would suggest that the relatives and carers believe day-to-day contact is vital for their reassurance.

4.8.4. Consideration for electricity monitoring

As highlighted in section 4.8.2, results from the survey have identified a number of considerations for the analysis of the electricity consumption data. The first of these was the need to monitor a range of different activities rather than just one activity. Therefore, for the subsequent monitoring of electricity usage, it was decided that the list of appliances to be monitored should include a range of appliances that represent different activities.

The second consideration for the analysis of the electricity consumption data was that, although the results from the survey highlighted the need to monitor a range of activities, there were some activities that were deemed to be more important than others, for example, food-related activities. Therefore, it was decided that the list of appliances to be monitored should highlight the importance of monitoring food and drink related appliances.

The list of activities that were chosen to be monitored based on the results from the survey aimed to highlight these two considerations, with a number of appliances chosen relating to food and drink preparation for example, the kettle, oven, electric hob, microwave, toaster and dishwasher as well as other appliances, for example, the electric shower, washing machine and television, which were also included to show a range of different tasks/activities.

The responses to the survey, as highlighted in section 4.8.3, also showed a number of considerations with respect to the collection, analysis and visualisation of the electricity consumption data. The first of these considerations was the need for the collection of the electricity consumption data to be undertaken in a non-intrusive manner. The use of an electricity monitor and data logger, described in section 3.3.2.1, addressed this consideration, as they could be considered non-intrusive in the way they monitor. The second consideration was the need for the results provided from the analysis to be reliable and accurate and to be presented in a way that was easily interpretable. This is because a relative/carer could not be expected to use the system, for example, if the developed model/system needed significant training for the results to be interpreted (Chan et al., 2008). In addition, for carers, especially paid carers who may have a number of clients, any significant variations in electricity usage would need to be clear, so that they were not overlooked. Equally, the system should not produce a high number of false negative alerts. Therefore, the method used to analyse the electricity consumption data must provide accurate and reliable recognition of each appliance, so as to be able to provide an accurate picture of appliances usage and to be able to see any important variations in the resident's appliance usage.

4.9. Conclusion

This chapter has presented the analysis of the results from the survey using three forms of analysis; statistical analysis, content analysis and thematic analysis. This chapter has also presented a discussion of the results from the survey, with the considerations used to inform the development of the study described in the following chapters. The next section of this thesis (Chapter 5-6) will present the analysis of the second set of data collected for this research, the whole house electricity consumption data.

Chapter 5: Trial Data Analysis

5.1. Introduction

As highlighted in Chapter 1, the data collection that was undertaken for this thesis was conducted in two parts. The preceding chapter (Chapter 4) has presented the results from the survey data collection, with the discussion in section 4.8, outlining the importance the respondents attached to different activities and tasks. This in part, informed the list of appliances to be monitored in the collection of electricity consumption data. As described in Chapter 3, whole house electricity consumption data and appliance usage diaries were collected from a number of households. This chapter will give an overview of the steps conducted to analyse the collected data to construct an initial model to recognise when an appliance (from the list in section 5.3) was used.

This chapter is divided into a number of sections. Section 5.2 gives a description of the method of data collection, including how the recorded electricity consumption data were processed into a form that could be used for the analysis. Section 5.3 of this chapter gives some an overview of the different structures and patterns of the electrical appliances, which were to be recognised. Section 5.4 provides a description of the process for the first iterations of the trial analysis of the electricity data, with a description of how the diary data were extracted, how the test and training data were developed and the construction of the recognition model with a discussion of how the results were interpreted. Section 5.5 gives a description of the subsequent iterations of the trial analysis, based on the issues discovered from the first iterations of the analysis. Finally, section 5.6 provides a summary of the analysis conducted, with an overview of the final model developed as well as a discussion of the issues discovered from the undertaking of this analysis.

5.2. Methods

For this part of the study the electricity consumption data and usage diaries were collected using the equipment as described in section 3.3.2. The equipment used was the mains sensor (as described in section 3.3.2.1.1), the electricity monitor (as described in section 3.3.2.1.2) and finally the data logger (as described in section 3.3.2.1.3). The participants were also asked to complete usage diaries based on the appliances that they had used during the collection period, from the list in section 3.3.2.2.

The participants of this study were approached based on a working or personal acquaintance with the researcher and invited to participate. The chosen participants had to meet the equipment requirements, as described in section 3.3.2, as well as have access to the Internet. Internet is required so that the data collected and stored by the data logger was periodically downloaded to a server. Section 5.2.2 describes how the data was accessed and downloaded from the server.

The electricity consumption data and diary data was collected for a one week period, between the 22^{nd} – 28^{th} November 2013.

5.2.1. Electricity and diary data collection - ethics

The collection of whole house electricity consumption data and diary data received ethics approval from the Information School Research Ethics Committee. A copy of the ethics approval and the information sheet are shown in Appendix 3.

5.2.2. The electricity data: pre-processing

As described in Chapter 3 (section 3.3), the data collection for this part of the project had two parts: the first part was the electricity consumption data of a household and the second was the diary data of usage of appliances whilst the electricity data was being collected. To undertake the analysis of the electrical consumption data it was necessary to reformat the collected data. As described in Chapter 3, the data were downloaded from the server as a series of zipped text files (shown as (.gz) in figure 5.1) and within each of the text files the data was stored as shown in Figure 5.1.

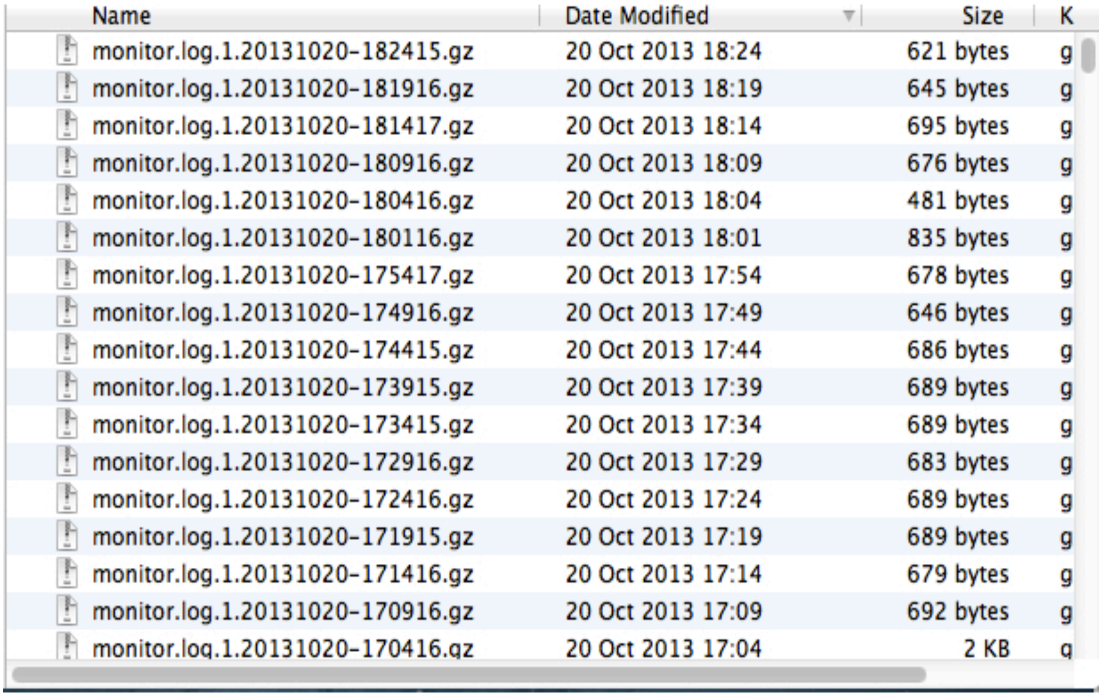

Name	Date Modified	Size	K
monitor.log.1.20131020-182415.gz	20 Oct 2013 18:24	621 bytes	g
monitor.log.1.20131020-181916.gz	20 Oct 2013 18:19	645 bytes	g
monitor.log.1.20131020-181417.gz	20 Oct 2013 18:14	695 bytes	g
monitor.log.1.20131020-180916.gz	20 Oct 2013 18:09	676 bytes	g
monitor.log.1.20131020-180416.gz	20 Oct 2013 18:04	481 bytes	g
monitor.log.1.20131020-180116.gz	20 Oct 2013 18:01	835 bytes	g
monitor.log.1.20131020-175417.gz	20 Oct 2013 17:54	678 bytes	g
monitor.log.1.20131020-174916.gz	20 Oct 2013 17:49	646 bytes	g
monitor.log.1.20131020-174415.gz	20 Oct 2013 17:44	686 bytes	g
monitor.log.1.20131020-173915.gz	20 Oct 2013 17:39	689 bytes	g
monitor.log.1.20131020-173415.gz	20 Oct 2013 17:34	689 bytes	g
monitor.log.1.20131020-172916.gz	20 Oct 2013 17:29	683 bytes	g
monitor.log.1.20131020-172416.gz	20 Oct 2013 17:24	689 bytes	g
monitor.log.1.20131020-171915.gz	20 Oct 2013 17:19	689 bytes	g
monitor.log.1.20131020-171416.gz	20 Oct 2013 17:14	679 bytes	g
monitor.log.1.20131020-170916.gz	20 Oct 2013 17:09	692 bytes	g
monitor.log.1.20131020-170416.gz	20 Oct 2013 17:04	2 KB	g

Figure 5.1: How the data are stored in a series of .gz files

time>1381667016</time><msg><src>CC128-v1.48</src><dsb>65535</dsb><time>12:12:50</time><tmpr>17.8</tmpr><sensor>0</sensor><id>03762</id><type>1</type><ch1><watts>00684</watts></ch1></msg>

Figure 5.2: Example of data recorded in text file

Each line within the zipped files (as shown in Figure 5.1) represents a new data point. An example of the data within each file is shown in Figure 5.2. Each line of the data file includes the time in UNIX time (in this example, 1381667016) and the associated power consumption (684W) at that point in time. The data also provided the current house temperature (i.e., 17.8C), the sensor number from which the data was recorded (0) and its ID number (03762). For this project, only one sensor was placed in each house. The sensor number, ID number and the recording of ambient temperature were not used for this research; however, the ambient temperature and additional in house sensors could be used for further research, if necessary.

To extract the data from these files, a java script was written that unzipped and combined every file into a single text file. This combined text file was then used for the subsequent analysis.

As discussed in Chapter 3, the analysis of the electricity consumption data was undertaken using the program Matlab. Matlab has a series of built-in functions that can be used to load data into its workspace (as described in section 3.4.5). The function 'textread[4]' was used to read each line of the text file and load the data into the workspace as a set of user-defined variables. Examples of these for this study include a variable that contains all the electricity consumption readings and another variable that contains the corresponding time recordings at which each electricity consumption reading was made. These variables were then visualised in graphical format using Matlab as shown in figure 5.3.

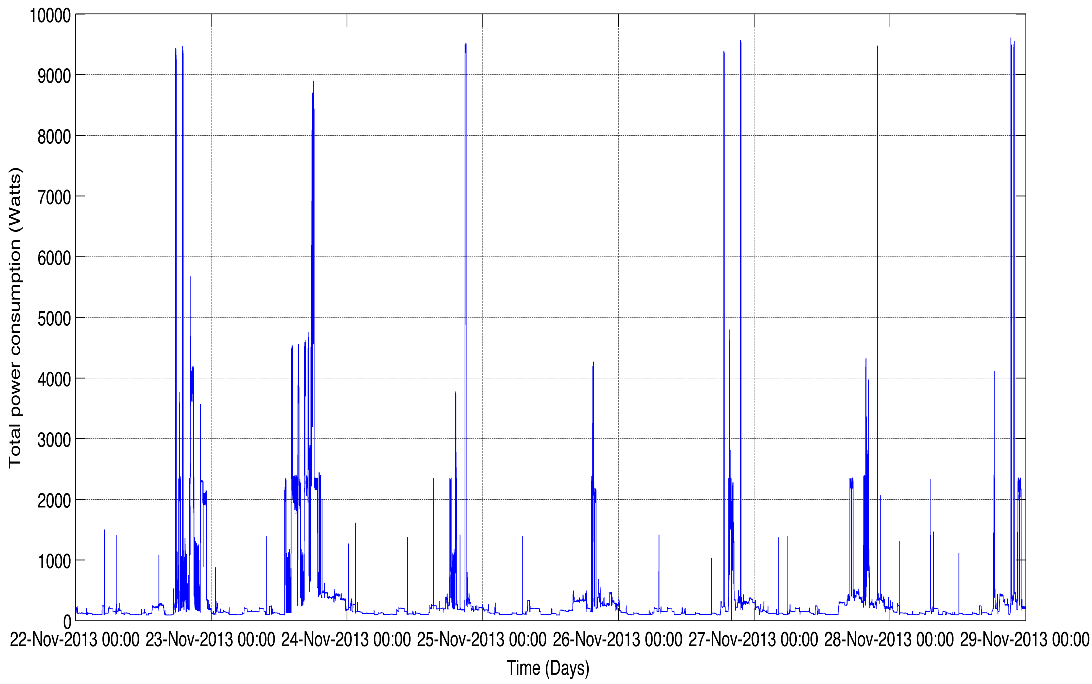

Figure 5.3: 7-day electricity consumption of a household using data extracted from the .gz files.

Figure 5.3 shows the total 7-day electricity consumption of the household used for the trial analysis. From figure 5.3, it can be seen that the electricity consumption of the household changed throughout the days: this corresponds with when different electrical appliances were turned on and off. By recording and storing the electricity consumption every six seconds, the monitors provided greater granularity that helped to demonstrate the different patterns of electricity usage for each appliance and it is these differences in electrical consumption that were analysed in the

[4] http://www.mathworks.co.uk/help/matlab/ref/textread.html

building of the recognition model to show when different appliances had been turned on and off. The following section (5.3) will show the differences between the electrical appliances and how these differences were used to develop a recognition model.

5.3. The electrical appliances

As shown in figure 5.3, the electrical consumption of a household changes as different electrical appliances are turned on and off. It is these differences that make the monitoring and recognition of electrical appliances via an electricity monitor possible, because different electrical appliances will show different patterns of electricity consumption. In some of the analyses of consumption data it was noted that some different appliances did have similarities in their consumption patterns. These similarities, and differences, and their impact on the development of a recognition model, will be discussed in section 5.3.1.

The list of the different electrical appliances that were switched on or off were recorded in the diaries and it is these appliances that the model was then developed to recognise their specific electrical signature from within the total electricity consumption data of the whole house. The appliances used were:

- Kettle
- Electric oven
- Electric hobs
- Television
- Washing machine
- Dishwasher
- Toaster
- Electric shower
- Microwave

As was discussed in Chapter 3, not all of the households that took part in this project had all of these appliances. For those houses, which recorded different appliances from the list, because of the use of gas (as outlined in section 3.3.2.2) more details will be given in the respective analysis sections.

The next section, (5.3.1) will discuss in detail the different electrical appliances and how their electricity consumption data were analysed, as part of the preliminary analysis of the electrical consumption data.

5.3.1. Electrical appliance signatures

The electrical household consumption data used for this analysis was from a household that had all of the electrical appliances on the list in section 5.3. The electrical household consumption data were recorded for one week from the 22nd to the 28th November 2013. Using the diary data and by processing the electrical consumption data, as described in section 5.2 and 5.3, the patterns of the electrical appliances were extracted and analysed, as discussed below.

5.3.1.1. The kettle

Electric kettles have a relatively simple electricity consumption pattern, in comparison with a number of other appliances. The user switches the kettle on; the electricity will heat up the element (which heats the water) and then will switch off when the water has reached the required temperature, either automatically at 100°C or by the user. The time that the kettle is on will be dependent on a number of factors, such as how much water is in the kettle, the temperature of the water when the switch is turned on, the required temperature of the water when the kettle will switch off. A typical example of the electricity consumption associated with this is shown in figure 5.4, this represents the electrical 'signature' of the kettle.

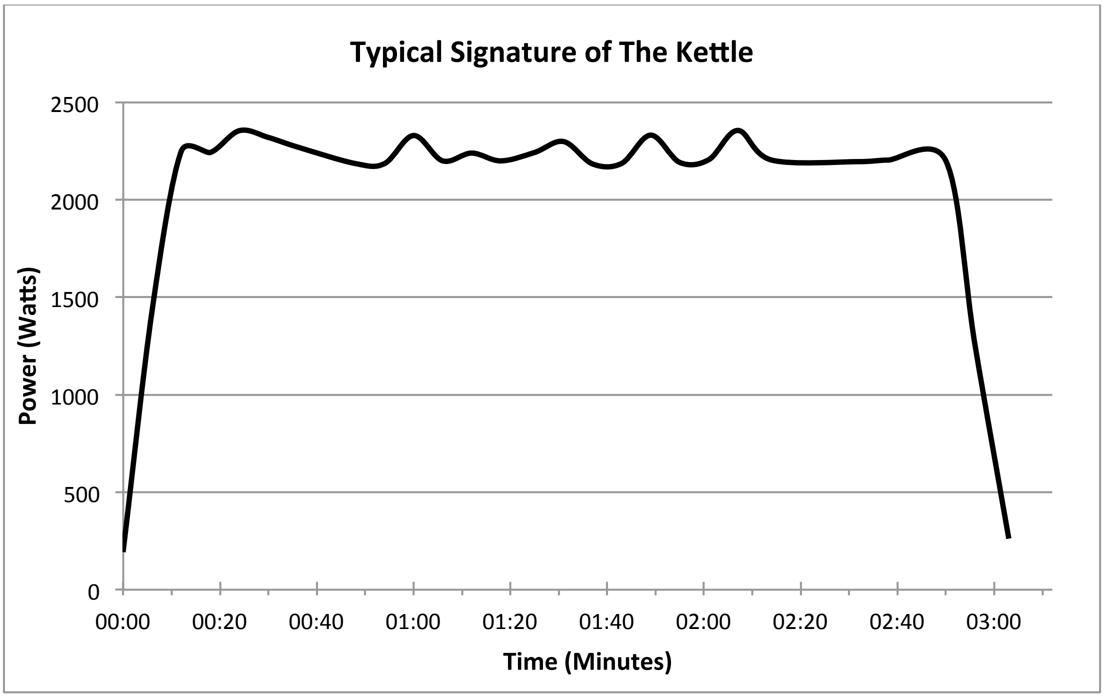

Figure 5.4: An example of the typical signature of the kettle shown graphically

5.3.1.2. The oven

The signature of the oven, as shown in figure 5.5, has a more complex pattern than illustrated for the kettle in Figure 5.4. The oven has an initial warm-up period in which electricity is used to heat the oven element to warm the air in the oven to the required temperature; once this is reached, the thermostat reduces the electricity use, and then will switch the element on and off and thus draw power at regular intervals to maintain the required temperature for as long as the oven is kept on by the user. The oven signature can vary according to the length of time the oven has been on, and by what setting the oven has been put on (for example convection oven, fan oven). The oven temperature program can also be altered by the operator by raising or lowering the required oven temperature.

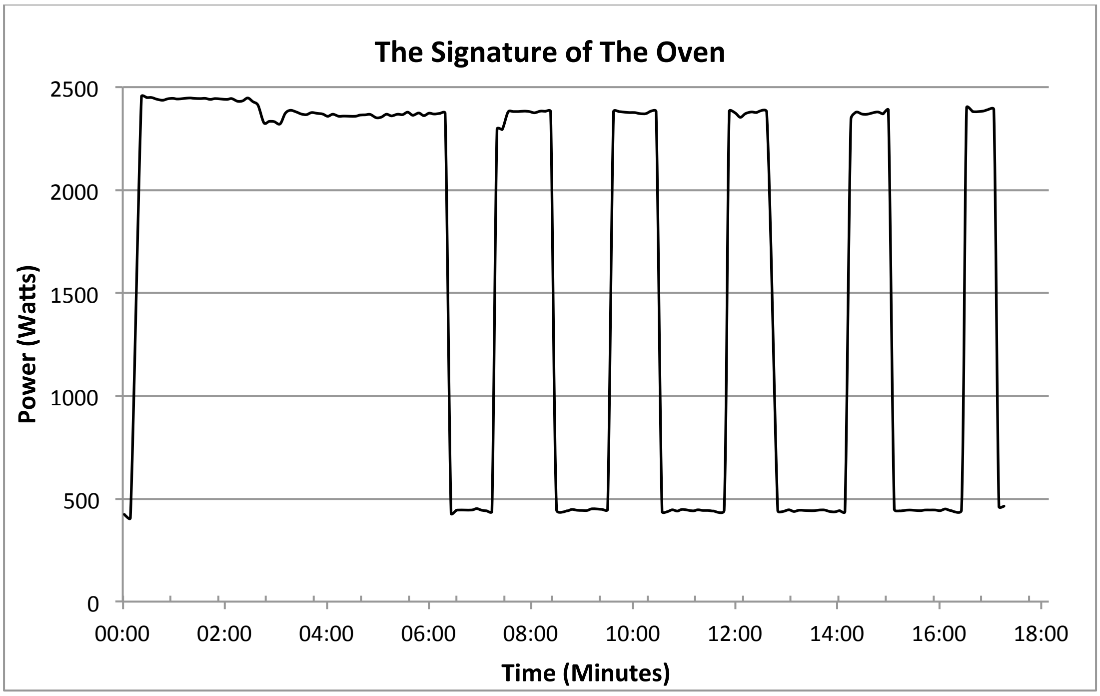

Figure 5.5: An example of the typical signature of the oven

For the preliminary analysis of the electricity consumption data, only the oven setting was used (a number of ovens have a grill facility but no grill signature data were recorded for this particular analysis).

5.3.1.3. The electric hob

The signature of the electric hob is shown in figure 5.6 and is very similar to that of the oven. The hob though was a more complex appliance to recognise because it has a multiple number of variable elements, and so can have many more variables that can affect the electricity consumption. The signature of the hob can depend on how long the hob has been put on for, so the length of the signature will vary, the number of hobs that have been turned on can vary and the different power settings that can be used for each of the hobs will also vary its signature.

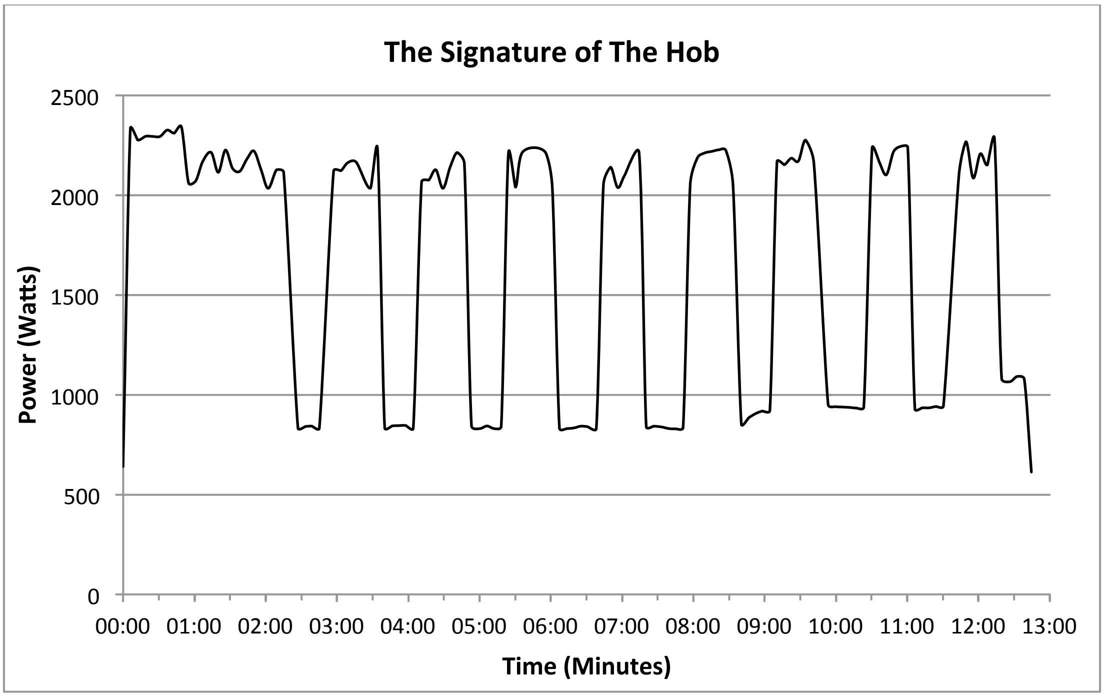

Figure 5.6: A typical example of the signature of the hob

For this preliminary analysis, only the times at which the hobs were turned on and off were recorded in the diaries, not the number of hobs that were used or the power settings at which they were set.

5.3.1.4. The television

The television is a very complex appliance to try and recognise because of the range of televisions available, e.g., because the age of the television and what type of TV it is (for example flat screen, LED, plasma etc.) affects the power drawn. Also the TV can be used with other appliances for example, digital versatile disk (DVD) players, games consoles, satellite boxes which further makes finding a unique electrical signature more difficult to identify.

The TV used in the preliminary analysis had a very low power draw, i.e., <100 Watts, that is almost impossible to recognise when it is turned on or off, against the noise generated whilst much higher power appliances, for example, the oven, are drawing power. For this reason, the television for this house was excluded from the analyses.

5.3.1.5. The washing machine

The signature of the washing machine is shown in figure 5.7. As figure 5.7 shows, the washing machine is another very complex pattern with different signatures for the different parts of the washing cycle, for example, heating the water, washing or spinning the clothes. A washing machine comes with many different features that the user can change, for example, the program setting, the temperature and the spin speeds, all of these will have an effect on the signature of the washing machine.

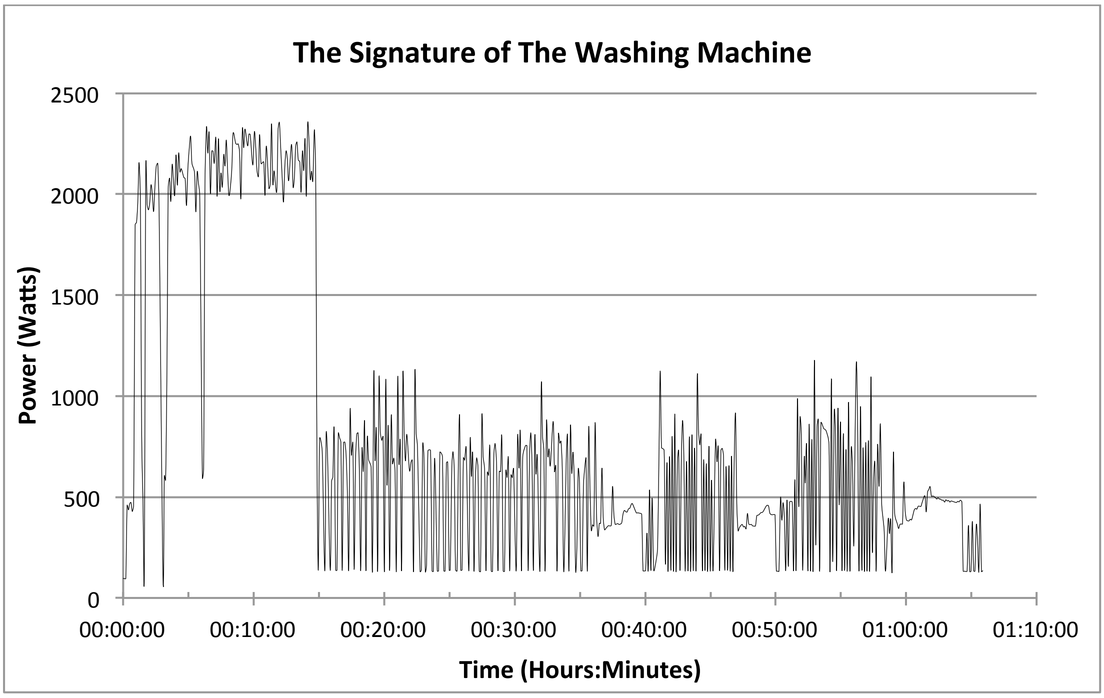

Figure 5.7: An example of the signature for the washing machine

For this preliminary analysis in the household diary, it was only recorded by participants that the washing machine had been used, not the settings used. It was also noted by the participant that the off time for the washing machine was not recorded in the diary, as they would leave the appliances to run, and then return after it was finished.

5.3.1.6. The dishwasher

The signature for the dishwasher is shown in figure 5.8. As figure 5.8 shows, the dishwasher draws power for a set period of time; it then draws minimal power for a set period of time and then repeats the electricity consumption pattern. As with the washing machine, a dishwasher can have different programs and times that can affect its signature, although for this preliminary analysis the differences were not shown, because all the recorded dishwasher signals were very similar. It was therefore concluded that the same program was used for all of the instances that the dishwasher was used. This conclusion cannot be proven as the participants were only asked to record when the dishwasher was turned on and off not the settings that were used.

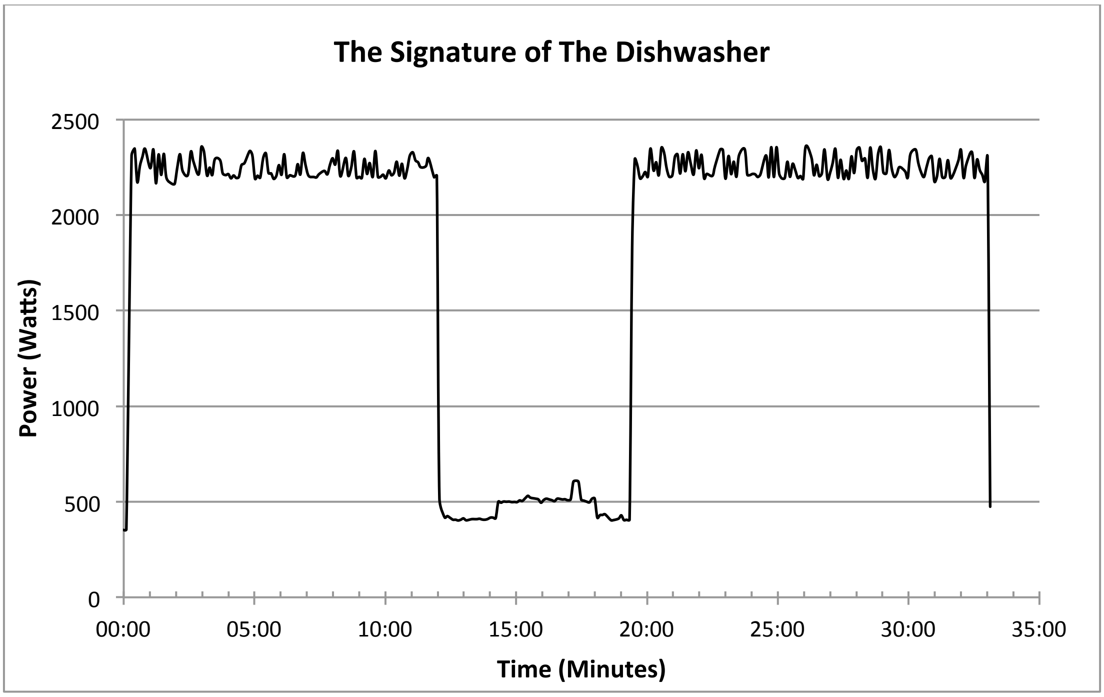

Figure 5.8: An example of the signature of the dishwasher

It was noted again by the participant, that for the preliminary analysis they did not record in the diary when the dishwasher completed its cycle and turned itself off. As with the washing machine, they said that they went away and left the appliance on and came back to it after it had finished.

5.3.1.7. The toaster

Figure 5.9 shows a typical signature for the toaster. The toaster has a very simple pattern, it is switched on, and it stays on for a required period of time and then turns itself off.

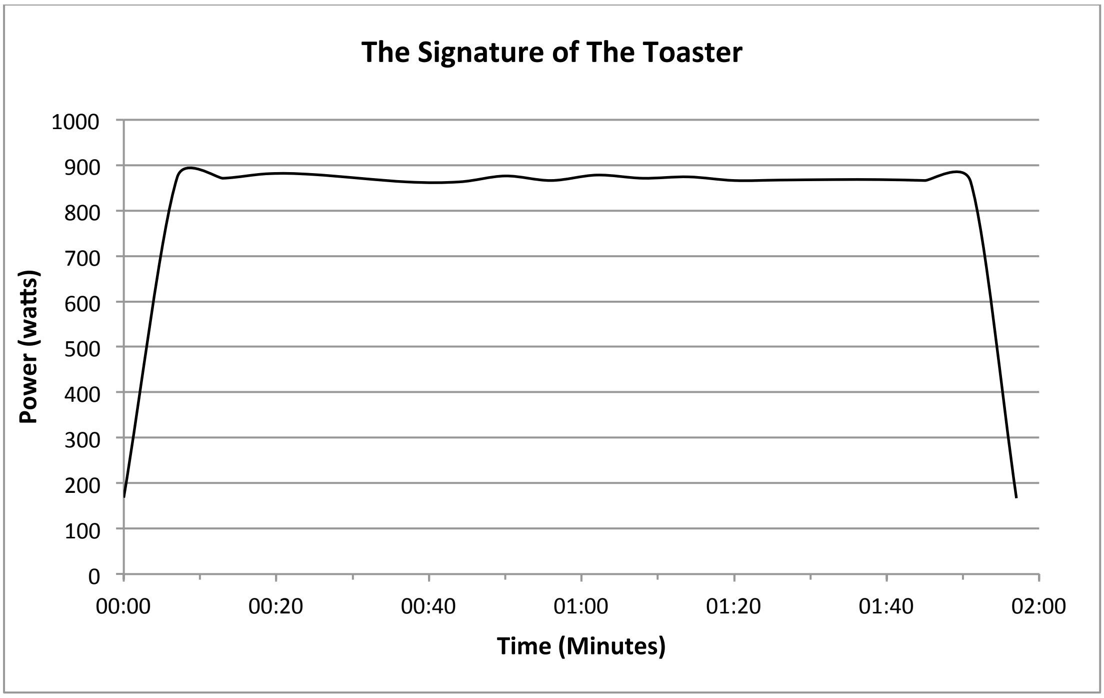

Figure 5.9: An example of the typical signature of the toaster

The toaster was not included in the preliminary analysis as there were not enough training and test instances; this is discussed in more detail in section 5.4.4. Other households that participated in this project used the toaster more often and the discussion of how it was analysed can be found in Chapter 6 (section 6.4).

5.3.1.8. The electric shower

The signature for the electric shower is shown in figure 5.10, and is a very distinctive, and therefore easily recognisable signature due to its high level of electricity consumption. For the shower, the maximum power used is approximately 9500 Watts, compared with only up to 2000 Watts for the appliances described previously. Although the power drawn by an electric shower can vary, depending on the make of the shower, from about 7.5 to about 10.5 Kilowatts, it is the biggest single power-drawing appliance in this project. The variation in the signature of the shower is the length of time the appliance is on, this will vary depending on how long the user uses the shower.

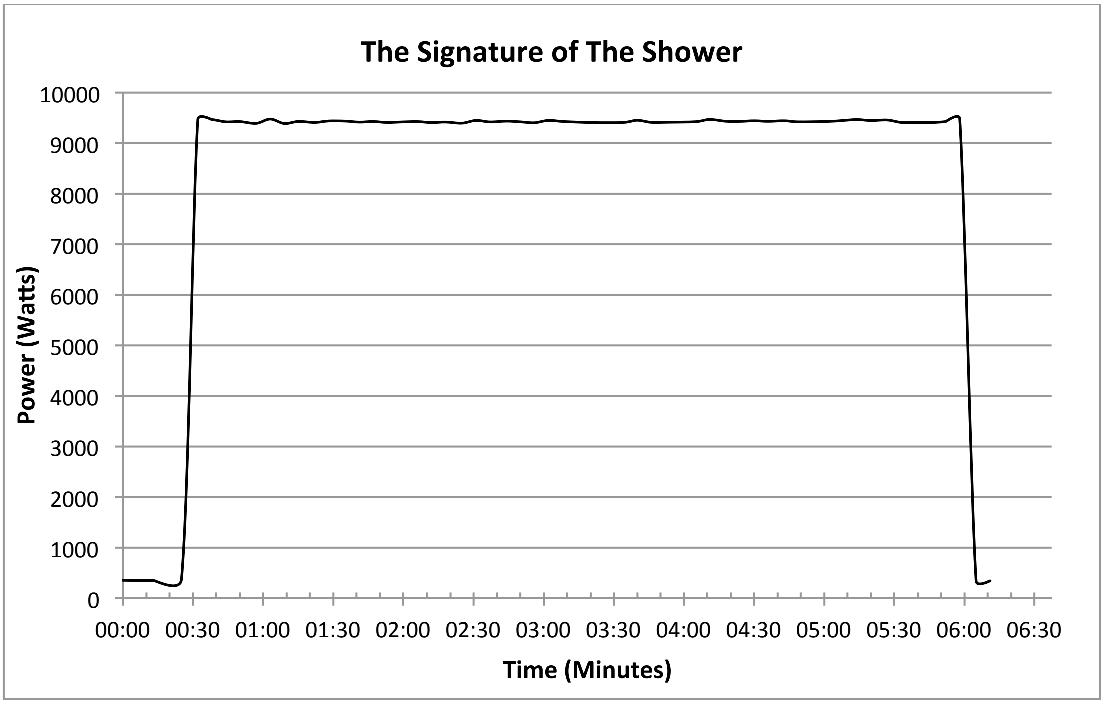

Figure 5.10: An example of the signature of the shower

5.3.1.9. The microwave

The signature of the microwave is shown in figure 5.11. As figure 5.11 shows the microwave draws power for a period of time and then turns off. The signature of a microwave will vary in the length of time that it is on for, this is a variable that is pre-set by the user and can be for a very short period of time (e.g., 10 seconds) or a longer period of time (e.g., 10 minutes). Other variations in the microwave signal can come from pre-set programs such as defrost or lower power cooking.

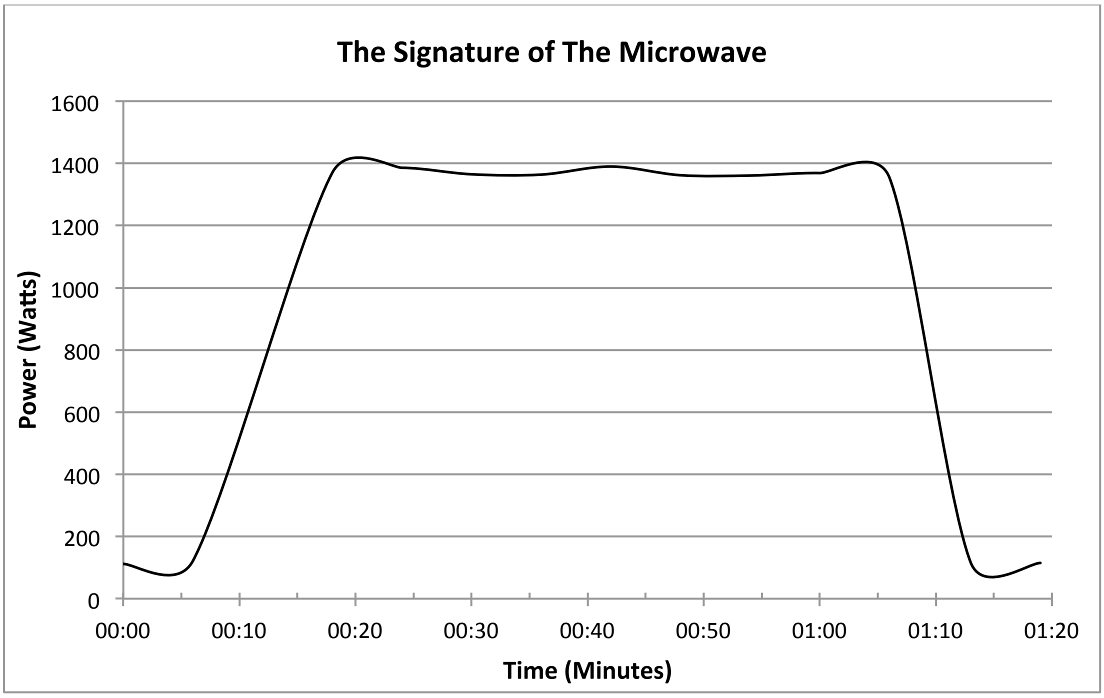

Figure 5.11: An example of the signature of the microwave

For the preliminary analysis, only the time at which the microwave was on for varied, not the signature size, so it was assumed that only one setting was being used. This could not be verified as the participants were only asked to record when the microwave was turned on and off and not the setting that was used.

5.3.2. Summary

As discussed and shown by all the different figures, different appliances have quite different electrical signatures, and different features, that can be used to identify the on and off times of an appliance from within the overall electrical consumption data. Section 5.4 will discuss how the different appliance electrical patterns were extracted and used to build a recognition model for these electrical appliances.

5.4. Appliance recognition

This section will give an overview of the methods used for the construction of the appliance recognition model used for the analysis of the electricity data. The section will give an overview of how the data were manipulated for this analysis and also how the training and test dataset were constructed.

5.4.1. Window design

The initial method used to transform the data for the appliance recognition model followed the methodology used by Lee et al., (2010) and Lin et al., (2010) described in the methodology chapter, section 3.4.1.

Figure 5.12 gives a tabular description of how a sliding window was used to transform the data set into the sliding window data set. For the sliding window in figure 5.12 the window size (i.e. the number of data points within the window) is 6 data points and the window shifts by 1 sample. Both the window size and the window shift can be varied; section 5.4.2 gives a description of how these are varied to form a key part of the analysis for this part of the project.

5.4.2. Trial 1- window design

As stated in section 5.4.1, this project followed the methodology developed by Lee et al., (2010) and Lin et al., (2010) to create a feature set by creating a sliding window of the data, as discussed in section 3.4.1.1. For this first trial of the window design, the window size adopted was six, i.e., six consecutive readings of electricity data, and the window shift used was 1. The window size of six was chosen as it equates to 36-second time slice and therefore close to the window size used for the method of Lee et al., (2010) and Lin et al., (2010). The features used for this first trial were also based on the features used by Lee et al., (2010) and Lin et al., (2010). The feature values used were:

- The raw data given from the electricity monitor, i.e., $(D_t, D_{t+1}, \dots, D_{t+(W-1)})$, where W is the window size, in this case 6 and D_t is the total electricity consumption for time t. The raw data are only included for reference and were not used for the training and testing of the recognition model.
- The mean value of the data points within each window, i.e., $D_{avg} = \left(\frac{(D_t, D_{t+1}, \dots, D_{t+(W-1)})}{W}\right)$.
- The peak value of the data points within each window, i.e.,

  $$D_{peak} = max\left(D_t, D_{t+1}, \dots, D_{t+(W-1)}\right).$$

- The standard deviation of the data points within each window, i.e.,

$D_{std} = \sqrt{\frac{\sum(x - \bar{x})^2}{W}}$ Where x is each value in the sliding window and $\bar{x}$ is the average of all of the data points in the sliding window.

- The root mean square of the data points within each window, i.e.,

$D_{rms} = \sqrt{\frac{D_t^2 + D_{t+1}^2 + \cdots D_{t+(W-1)}^2}{W}}$.

- The peak to average ratio, i.e., $D_{ptavg} = \frac{D_{peak}}{D_{avg}}$
- The peak to root mean square ratio, i.e., $D_{ptrms} = \frac{D_{peak}}{D_{rms}}$
- The root mean square to average ratio, i.e., $D_{rmstavg} = \frac{D_{rms}}{D_{avg}}$

Using the features given in the list above the electricity consumption data were transformed from the original electricity consumption data into the feature data set, an example of this is shown in figure 5.12.

Electricity consumption data
221
213
219
214
217
215
212
214

Window number	Raw window data						Average	Peak	Standard deviation	Root mean square	Peak to average ratio	Peak to root mean square ratio	Root mean square to average ratio
1	221	213	219	214	217	215	216.5	221	3.0822	216.5183	1.0208	1.0207	1.000085
2	213	219	214	217	215	212	215	219	2.6077	215.0132	1.0186	1.0185	1.000061
3	219	214	217	215	212	214	215.17	219	2.4833	215.1786	1.0178	1.0178	1.000055

Figure 5.12: Transformation of the electricity data (measured in Watts) to the window data and feature set – this figure shows three rows of consecutive sliding windows of window size 6. The colours indicate the exact place in the window number sequence for each reading for clarity.

5.4.3. Trial 1- diary data

The appliance usage data were recorded in diaries as described in Chapter 3 (section 3.3.2.2). For the next part of the analysis the data from these diaries had to be extracted and the data split into training and test sets.

For this first trial, only four of the recorded electrical appliances, from the list in section 5.3, were used for the recognitions. These appliances were:

- The electric shower, which is a very high-powered appliance.
- The microwave, which is a low powered appliance.
- The kettle, which is a medium powered appliance.
- The dishwasher, which is a medium powered appliance.

To explain how the diary data were used to create a test and training set of data this section describes a worked example for the electric shower.

Within this trial period, the electric shower was used eight times by the participant and the dates of use, as well as the times the appliance was switched on and off, were recorded by the participants using the diaries.

To match the time and date given on the diary to the corresponding feature data a graphical user interface (GUI) was created, as shown in figure 5.13.

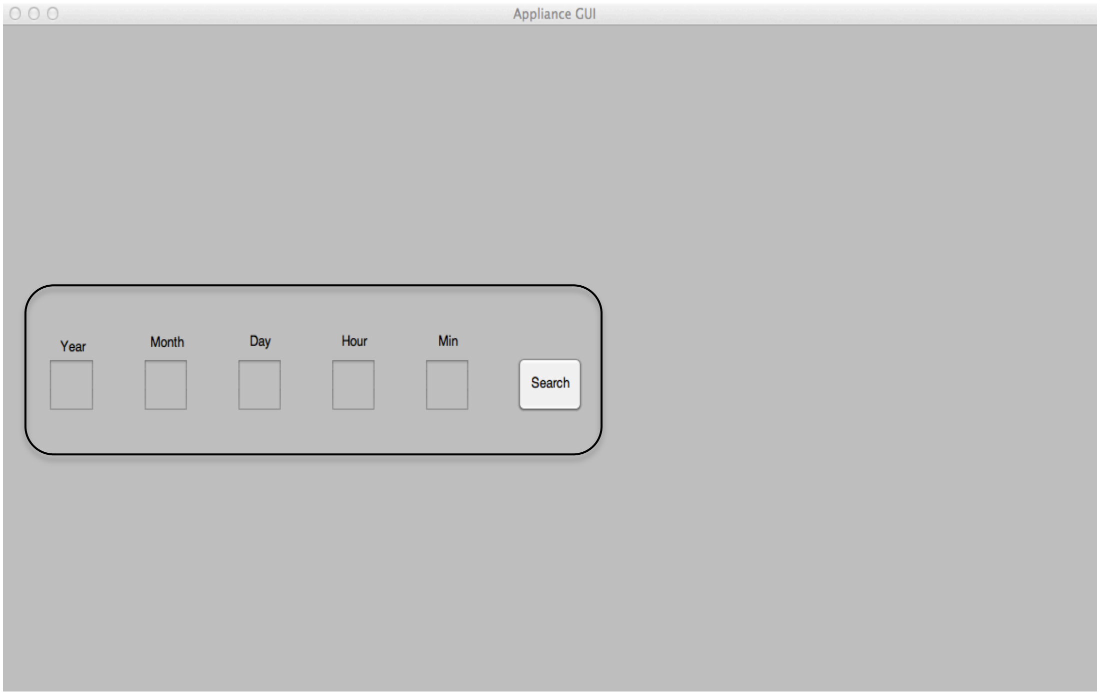

Figure 5.13: GUI used for finding appliance data

For the GUI in figure 5.13, the approximate date and time for the use of one instance of the shower (as recorded by the participant in the diary) was entered into the corresponding boxes and the search button was pressed. The search function identified this time and retrieved the corresponding data in the feature set for the entered time and date (as well as the data for the previous and following two minutes) then produced the table and graph as shown in figure 5.14.

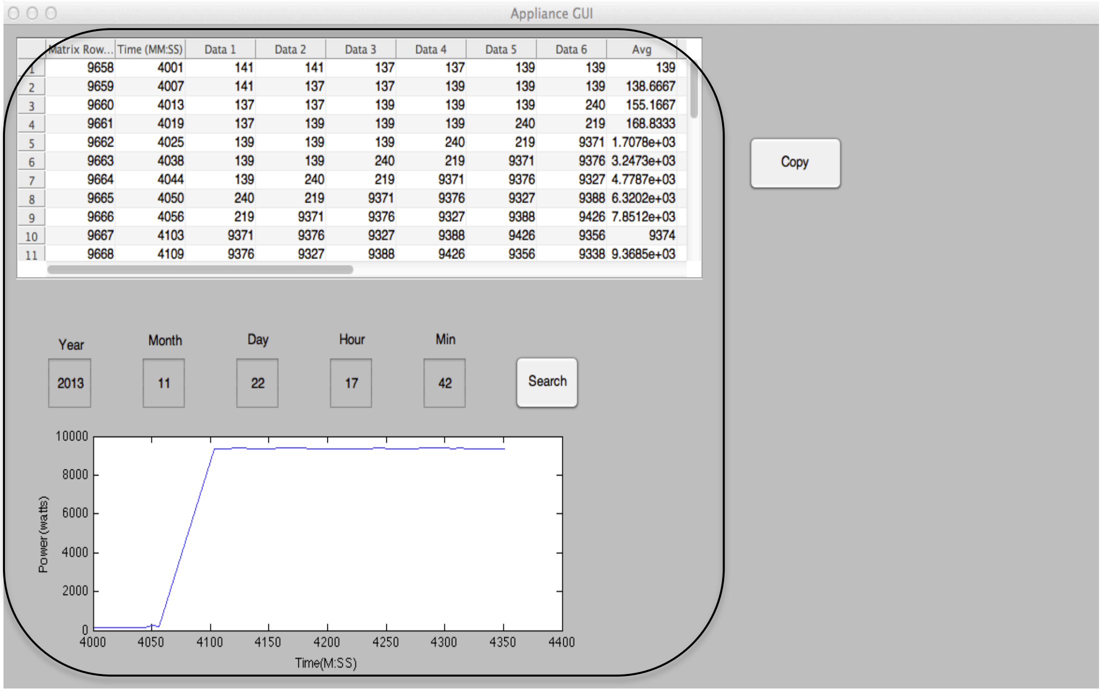

	Matrix Row...	Time (MM:SS)	Data 1	Data 2	Data 3	Data 4	Data 5	Data 6	Avg
1	9658	4001	141	141	137	137	139	139	139
2	9659	4007	141	137	137	139	139	139	138.6667
3	9660	4013	137	137	139	139	139	240	155.1667
4	9661	4019	137	139	139	139	240	219	168.8333
5	9662	4025	139	139	139	240	219	9371	1.7078e+03
6	9663	4038	139	139	240	219	9371	9376	3.2473e+03
7	9664	4044	139	240	219	9371	9376	9327	4.7787e+03
8	9665	4050	240	219	9371	9376	9327	9388	6.3202e+03
9	9666	4056	219	9371	9376	9327	9388	9426	7.8512e+03
10	9667	4103	9371	9376	9327	9388	9426	9356	9374
11	9668	4109	9376	9327	9388	9426	9356	9338	9.3685e+03

Figure 5.14: GUI showing the data and graph for the corresponding time and data

The researcher was then able to select the row of the table that represented when the appliance had been switched on; this is shown in the black circle in figure 5.15.

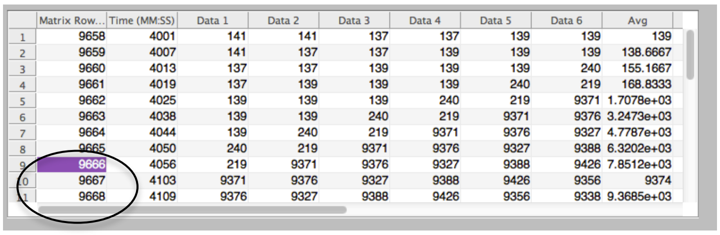

	Matrix Row...	Time (MM:SS)	Data 1	Data 2	Data 3	Data 4	Data 5	Data 6	Avg
1	9658	4001	141	141	137	137	139	139	139
2	9659	4007	141	137	137	139	139	139	138.6667
3	9660	4013	137	137	139	139	139	240	155.1667
4	9661	4019	137	139	139	139	240	219	168.8333
5	9662	4025	139	139	139	240	219	9371	1.7078e+03
6	9663	4038	139	139	240	219	9371	9376	3.2473e+03
7	9664	4044	139	240	219	9371	9376	9327	4.7787e+03
8	9665	4050	240	219	9371	9376	9327	9388	6.3202e+03
9	9666	4056	219	9371	9376	9327	9388	9426	7.8512e+03
10	9667	4103	9371	9376	9327	9388	9426	9356	9374
11	9668	4109	9376	9327	9388	9426	9356	9338	9.3685e+03

Figure 5.15: GUI showing the selection of the data from the table

The selected data row was then copied into a new table, using the copy button (as shown in figure 5.16).

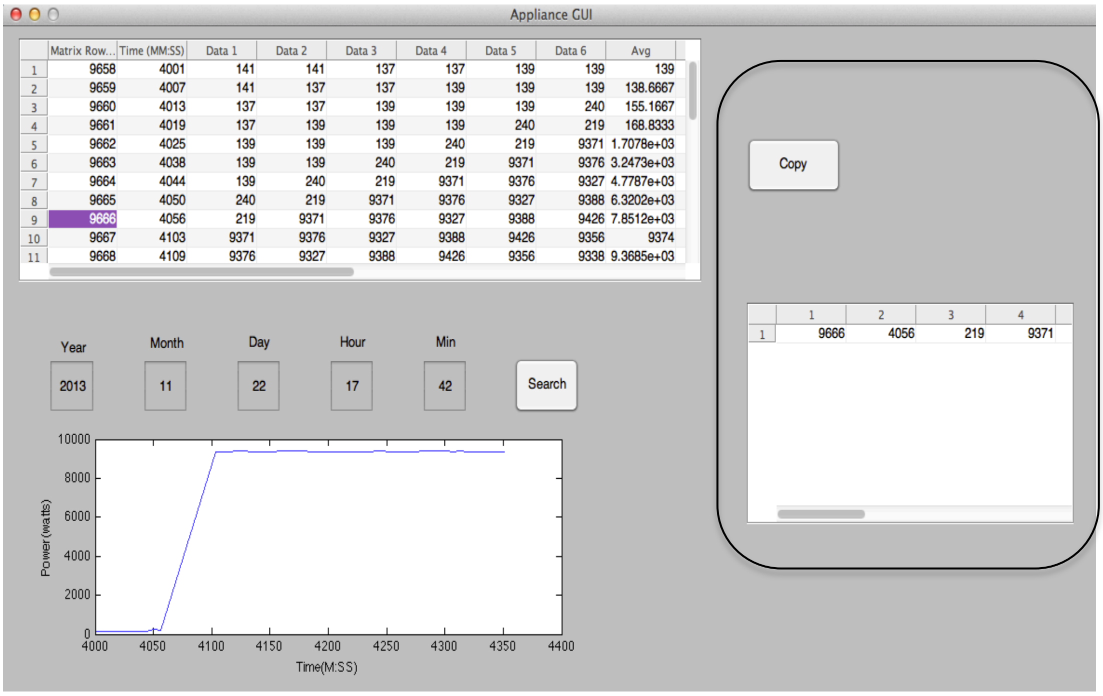

Figure 5.16: GUI used for selecting appliance data

This process was then repeated for the next time the shower had been used and repeated until all the times that the shower had been used had been entered. Once all of the diary data had been entered into the GUI and the results copied into the new table, the table was then saved as a Matlab MAT-file.

This whole process was then repeated for all of the appliances that were recorded so that each appliance had a saved MAT-file of all the instances the appliance had been used. Each of these files were then used later in the analysis, where they were then split into a test and training dataset based on the dates used for training and testing as outlined in section 5.4.5.

5.4.4. Trial 1- development of training and test datasets

As discussed in the literature review in section 2.4.5, there are a number of different methods for choosing how to structure the data so as to provide training and test data. The choice of the structure of the data for training and test data for electricity data was constrained by the limited data, in terms of appliances usage and also the aim of this work to provide data in terms of days of usage. This project recorded a full week of one household's electricity consumption data, which provided a large amount of data points, although it also highlighted an unforeseen issue regarding a household's usage of certain appliances. For example, the dishwasher was only used four times throughout the week, which limited how data could be separate into training and test datasets. Due to the restriction of the limited size of the dataset, it was decided to use the hold-out method for splitting the data into two datasets, a training set and a test set, as described by Witten et al., (2011) and in section 2.4.5.

The use of the holdout method has the limitation that the model could overfit (Rogers & Girolami, 2012) as described in more detail in 2.4.5.8, as well as it possibly not being generalizable to the overall population (Witten et al., 2011). As discussed in section 2.4.5.9, a method for overcoming overfitting is to conduct cross-validation. The use of cross validation for the data in this thesis is explored in more detail in section 5.5.6.

For the hold-out method used for this trial, the data were split based on days, with the training days being the 23rd, 24th, 26th and 28th and the test days being the 22nd, 25th and 27th November (2013). The days for training and testing had to be carefully selected by the researcher in order to provide a minimum of two data points for the training set, as required by the recognition method. The distribution of the appliances in terms of training and test dataset is shown in table 5.1. It can be seen from table 5.1 that, in the trial household, some of the appliances were used a limited number of times, this restricted the number of instances that could be allocated to the training and test data sets.

Appliance class membership	Training instances	Test instances	Total
Dishwasher	2	2	4
Electric hob	4	2	6
Kettle	3	2	5
Microwave	6	4	10
Oven	2	2	4
Shower	5	3	8
Toaster	1	1	2
Television	6	3	9
Washing machine	5	2	7

Table 5.1: Distribution of appliances in the training and test datasets for the trial house

5.4.5. Trial 1- recognition model

To code the first trial of the recognition model a Matlab function was created that transformed the electricity consumption data into the feature set as discussed in section 5.4.2. The feature set was then spilt into a training and a test set, the training data for this trial was four complete days of electricity usage with the test data being three complete days, as discussed in section 5.4.4.

As discussed in section 3.4.3, the use of a classification method was chosen as the method for conducting the recognition of the appliances from the electricity data. This was instead of the other approaches discussed in sections 2.4.5 and 3.4. This is due to the structure of whole house electricity data, i.e., the whole house electricity data represents a cumulative value of the electricity consumption of the house and thus presents cumulative patterns of appliances being used at that time rather than individual appliance patterns. Secondly, the length of the pattern to be matched would also be highly variable; as described in section 5.3.1, the length of time that appliances would be used would vary, e.g., the amount of water in a kettle would be one of several factors affecting how long it took the water to boil, and therefore for the kettle to remain on. Therefore a classification method was seen as more suited to this type of analysis than the use of pattern matching. As discussed in section 3.4.3 the classification method chosen for this analysis was a naïve Bayes classifier.

To train the naïve Bayes classifier a class file was created, which related the features in the training data to which appliances had been turned on, based on the times given by the diary data as discussed in figure 5.4.3. Each line in the class file contained either ‘off’, ‘shower on’, ‘dishwasher on’, ‘kettle on’ or ‘microwave on’ depending on what, if anything, had been turned on at that time.

The training data and the class file were then used to train the inbuilt Matlab naïve Bayes classifier function[5]. Once the classifier had been trained, using the inbuilt Matlab function, it was tested using the test data to predict the appliances being ‘on’ or ‘off’. Table 5.3 shows the recognition results from these test data for each of the four appliances, for windows six, which the model was trained to recognise.

5.4.5.1. Trial 1- recognition model: random baseline

To create baseline results to provide a comparison for this trial, the class file, as discussed in section 5.4.5, was randomised for the ‘shower on’, ‘dishwasher on’, ‘kettle on’ or ‘microwave on’ data points. Thus, for this baseline, the class assigned to the point, which is, for example, ‘shower on’, was randomly chosen out of the four appliances. The points classed as ‘off’ were not changed for this baseline. As

[5] http://www.mathworks.co.uk/help/stats/naivebayes-class.html

discussed in section 5.4.5, the training data and the random baseline class file were used to train the classifier. This random baseline classifier was then tested using the test data. The results given from this random baseline are shown in table 5.2.

Appliance	True Positive	False Positive	False Negative	True Negative	Sensitivity	Specificity	Positive predictive value	Negative predictive value	Total accuracy
Shower	0	85	3	39198	0%	99.784%	0%	99.992%	99.776%
Microwave	0	202	4	39080	0%	99.486%	0%	99.990%	99.476%
Kettle	0	12	2	39272	0%	99.969%	0%	99.995%	99.964%
Dishwasher	0	4	2	39280	0%	99.990%	0%	99.995%	99.985%

Table 5.2: Results from the random baseline

5.4.6. Trial 1- recognition results and discussion

Appliance	True Positive	False Positive	False Negative	True Negative	Sensitivity	Specificity	Positive predictive value	Negative predictive value	Total accuracy
Shower	3	2	0	39281	100%	99.995%	60%	100%	99.995%
Microwave	3	3	1	39279	75%	99.992%	50 %	99.997%	99.990%
Kettle	1	384	1	38900	50%	99.023%	0.260%	100%	99.020%
Dishwasher	1	48	1	39236	50%	99.878%	2.041%	100%	99.875%

Table 5.3: Results from the first trial to recognise the appliances from the electricity consumption data

The results shown in table 5.3 shows that there were differences among the recognition rates for the four appliances. For example, the best recognition rate of all the appliances was for the shower, which had an overall accuracy of 99.995%, and this included only 2 false positives and no true negatives. In contrast, the worst recognition rate was for the kettle, which had an overall accuracy of 99.020%, with 348 false positives. The sections below outline some of the reasons for these differences in recognition performance between the appliances.

Comparing the results in table 5.3 to the random baseline results in table 5.2, it can be seen that the results from table 5.3 produce a better result both in terms of overall accuracy for the shower and the microwave but poorer results for the kettle and the dishwasher. Although, the random baseline does not correctly recognise any of the appliances being turned on, which is a requirement of the model. As discussed in more detail in section 5.4.6.6, the use of total accuracy as a measure of the performance can be seen as flawed due to the highly skewed dataset, caused by the limited use of some appliances over the test period. The use of the total accuracy as a measure of performance failed to show that the requirement to recognise at least once instance when each of the appliances was actually used, had been met. These two factors are the reasons why the random baseline gave better results, in terms of total accuracy for two appliances (the kettle and the dishwasher) compared with the results from trial 1 in table 5.3. However, when comparing the random baseline results, in terms of PPV and sensitivity, the random baseline gives a result of 0% for all appliances. This therefore highlights the flaw of using total accuracy as a measure of the performance of a model as well as highlighting why PPV and sensitivity are better measures of the performance, as discussed in more detail in section 5.4.6.6.

5.4.6.1. The shower

As can be seen from table 5.3, for the shower, the recognition model recognised all of the instances in which the shower was used, but also produced two false positives when the model predicted that the shower had been turned on when it had not.

The prediction given by the model for the shower were analysed further to understand why the model was predicting 2 false positives. This analysis found that the model was also predicting when the shower had been switched off not when it was just switched on. The reasons for this were found to be how the feature set was calculated from the windows. The data used for this model were based on a window size of six data points, as previously discussed in section 5.4.2. The data points selected for the training data were when the appliance was turned on, the training set for the shower is shown in table 5.4.

Window number	Raw Data (reference data only)						Average	Peak	Standard Deviation	Root mean square	Peak to average ratio	Peak to root mean square ratio	Root mean square to average ratio
37547	269	9372	9411	9499	9345	9426	7887	9499	3732.414	8591.499	1.204	1.106	1.089
62517	302	9338	9337	9337	9324	9327	7827.5	9338	3686.732	8520.353	1.193	1.096	1.089
64063	376	9572	9421	9440	9488	9444	7956.833	9572	3714.229	8649.130	1.203	1.107	1.087
90028	406	9610	9457	9432	9504	9486	7982.5	9610	3712.217	8672.03	1.204	1.108	1.086
90291	422	9430	9402	9435	9432	9394	7919.167	9435	3672.886	8599.706	1.191	1.097	1.086

Table 5.4: The shower training data set

The data points used for the training data sets were the points when the appliance was first turned on, this is shown by the big increase in power between the first and second columns of the raw data in table 5.4. When the appliance was turned off there was a similar large change, albeit a decrease, rather than an increase, in power between the fifth and sixth columns of the raw data, as shown by the shaded rows of the feature set from the recognition results of the shower in table 5.5. As the feature set used for this model was calculated from the window data, the calculation

is unaffected by the position of the change in power and its absolute value, i.e., whether it is a large increase or decrease. Therefore, the model recognised when the shower had been turned off, but identified it as being turned on, because the features were similar to those in the training set, even though there was a large decrease in power use, rather than an increase. This is shown by the feature set of the results from the recognition model for the shower in table 5.5. It was noted that for one of the times that the shower was turned off, it was not recognised by the model, this is because there were other appliances on at the same time that the shower was turned off, so the feature set values were different. The problem with the values of the feature set being distorted when other appliances are on in the background is discussed in further detail in section 5.4.6.3.

Window number	Raw Data (reference data only)						Average	Peak	Standard Deviation	Root mean square	Peak to average ratio	Peak to root mean square ratio	Root mean square to average ratio
9666	219	9371	9376	9327	9388	9426	7851.167	9426	3739.118	8561.055	1.201	1.101	1.090
9720	9229	9243	9222	9282	9228	217	7736.833	9282	3684.015	8436.146	1.200	1.100	1.090
10333	818	9389	9339	9293	9246	9196	7880.167	9389	3460.404	8489.740	1.191	1.106	1.077
77217	371	9465	9466	9421	9426	9391	7923.333	9466	3699.983	8613.213	1.195	1.099	1.087
77262	9409	9409	9408	9427	9478	341	7912	9478	3709.115	8606.065	1.198	1.101	1.088

Table 5.5: Showing the corresponding feature set for the results of the recognition model for the shower

This revealed a limitation for this particular window method, i.e., that the direction of the power change within the window does not affect the calculations, and therefore, as in the case for the shower being turned off, will give incorrect results, in this case false positive results.

5.4.6.2. The microwave

In the case of the microwave, the model recognised three of the four instances it was turned on. There were also 3 false positives and further analysis into the reasons for these identified the same problem as with the shower. The false positives given by the model were when the microwave was turned off and were due to the calculation of the feature set, as described in section 5.4.6.1.

5.4.6.3. The kettle

For the kettle, the model recognised one out of the two instances when the kettle was used but it also gave 384 false positive predictions. Further analysis into the reasons for these results led to the identification of two problems with this method, described in the following two paragraphs.

The first problem was that the power rating for the kettle in this household was found to be very similar in peak size to other appliances, e.g., the washing machine. The washing machine was not included in this first trial, so it is unclear if there would still be the large number of false positives if the model had been trained to recognise the washing machine as well.

The second problem identified with the analysis of the kettle results was that this method only gave good recognition results when the appliance being recognised was the only appliance on at that time. The reason for this is how the window data and feature sets were calculated, which is indicated in table 5.6.

Window number	Raw Data (reference data only)						Average	Peak	Standard Deviation	Root mean square	Peak to average ratio	Peak to root mean square ratio	Root mean square to average ratio
12063	849	2423	2413	2131	2417	2421	2109	2423	627.902	2185.505	1.149	1.109	1.036
76136	2228	4179	4180	4291	4293	4307	3913	4307	827.501	3985.248	1.101	1.081	1.018
88387	690	2334	2319	2393	2381	2400	2086	2400	684.760	2177.806	1.150	1.102	1.044

Table 5.6: Examples of some of the data points of the kettle

Table 5.6 shows some of the data points for when the kettle was switched on in the data but also highlights the problem found with the design of this method. As the feature set was calculated from the window data, if there were other appliances on when that appliance came on then the calculation for that feature point would be masked by the higher overall power of the raw data, as row 2 in table 5.6 shows. The calculations used for the window data and feature set made no allowances for when there were other appliances on at the same time. This was a problem that was addressed for the next test (section 5.4.7).

5.4.6.4. The dishwasher

In the case of the dishwasher, the model recognised one of the two instances when it was turned on. There were also 48 false positives which, following further analysis, were identified to be due to two reasons. The first reason was that the dishwasher in this household was very similar in power consumption to another appliance, the oven, that was not used in this test and that was the reason for some of the false positives. The second reason was as discussed in section 5.4.6.1, the model was also recognising when the dishwasher was turned off as well as on.

5.4.6.5. Conclusion

This trial and the window design showed a good recognition rate when no other appliance was on but could not recognise specific appliances when multiple appliances were on at the same time. The exception to this was the shower; this was because the shower had such a high power use relative to the other appliances in this house that it was easily recognised even if other appliances were already on.

The method also highlighted the problem with the calculations used for the feature set and how some of the results, especially those for the microwave, could have been influenced by the fact that the microwave was only used when there were no other appliances on. There was also a lot of misclassification among certain appliances that were not included in this preliminary analysis, but which had very similar power draws to the appliances that were included, namely the oven, dishwasher, kettle and washing machine.

This analysis also highlighted a limitation with the use of a sliding window within the window design method. This limitation was that if two appliances were to be turned on within the same window it would be impossible for the model to recognise either of those appliances. The reason for this is that the feature set is calculated from the window data and if there were two appliances turned on within the same window the feature set would not represent the training data for either of the appliances.

5.4.6.6. Next steps

The results from this trial as described and discussed in 5.4.6.1-5.4.6.5, highlighted a number of issues and problems with the methodology that need to be addressed for the next stage of the analysis.

The first problem highlighted was that the model reported an appliance was being turned on when it was actually being turned off, as well as when it was being turned on, leading to a number of false positive predictions. This was addressed in the next stage of the analysis by applying a post analysis filter, so that only the points where there was a power increase between the first and second point would be considered.

The second problem was how the feature set was calculated if there were other appliances on already, when an appliance the model was trying to recognise was turned on, as discussed with the kettle data in section 5.4.6.3. The next stage was therefore to outline ways that the calculations of the feature set could be redesigned so that the model had a better recognition rate for appliances that were turned on when other appliances were already drawing power.

The third problem highlighted was that of data misclassification, where the model identified appliances not included in this initial trial, for example, the oven and the washing machine, as one of the trial appliances that was included. The next step was to include all the appliances and then evaluate this to see whether the model still misclassified the data.

The fourth problem was how the overall results for the model were shown. For this original trial, the results from the model were given as total accuracy percentages. These total accuracy percentages were very high for all of the appliances even with the large number of false positives that were produced by the kettle. It was decided for all the future analyses to show the overall results for each window size as the overall positive predicted value (PPV) and the overall sensitivity, rather than the overall accuracy of the model. The reason for this is that the aim of this model is to recognise correctly when the appliance has been used, so as to indicate an activity by a person, and therefore that a person living on their own is well enough to be doing that activity. Therefore, the aim is to have a model with a high overall PPV percentage, i.e., so that of the instances when the model predicts that an appliance is being used, that it is actually being used, and someone is performing an activity and are, therefore, presumably well. This model also needs to be able to recognise, with a level of high accuracy, when an appliance has not been used, so as to indicate an activity has not been performed and a person is possibly unwell. To achieve this the model also need to have a high overall sensitivity.

The overall PPV and overall sensitivity are calculated to give the overall results of the model not individual appliances within the model and are therefore calculated as shown in the equations below.

$$Overall\ PPV = \frac{(TP_1 + TP_2 \cdots\ TP_n\ )}{\big((TP_1 + TP_2 \cdots\ TP_n\ ) + (FP_1 + FP_2 \cdots\ FP_n\ )\big)}$$

$$Overall\ sensitvity\ = \frac{(TP_1 + TP_2 \cdots\ TP_n\ )}{\big((TP_1 + TP_2 \cdots\ TP_n\ ) + (FN_1 + FN_2 \cdots\ FN_n\ )\big)}$$

Where TP, FP and FN represent true positive, false positive and false negative for each appliance and n is the number of appliances, which the model is trained to recognise.

Apart from the problems highlighted by the results from this first trial of the recognition model, this process has also highlighted other areas that needed to be investigated in the next stage of the analysis. The first area is changes in the window size. For the initial trial only the results from one window size (6 data points) were analysed. The next stage of this analysis should investigate if changing the window size might give a better appliance recognition rate.

The second area is the analysis of the attributes selected for the feature set. The attributes for this initial trial were taken from the method of Lee et al., (2010) and Lin et al., (2010). The next stage was to investigate whether changing the attributes within the feature set would also allow better appliance recognition rates. These are explained in the following section (5.4.7).

5.4.7. Trial 2

As discussed in the previous section, the first analysis of the electricity consumption data identified many problems and indicated potential areas for investigation, which this second trial analysis attempted to address.

The first area of investigation that was highlighted in section 5.4.6.6 was that the first trial analysis was only run for one window size (window size 6). For all future analyses of the electricity data it was decided that the model would be run for window sizes 3 (covering 18 seconds) to 10 (60 seconds). For reference, the Matlab script for the first trial results (section 5.4.5) was also re-run for window sizes 3 to 10 and the results table for this can be found in table A4.1 of Appendix four.

As described in section 5.4.6, another of the problems with the initial window design was that if there were appliances on in the background when the appliance that the model had to recognise came on, the model did not recognise it, as highlighted in section 5.4.6.3. The cause of this problem was the way the feature set was calculated from the window data. To address this problem the window data was redesigned so that the difference between the values in the window data was calculated as, for example: $D_{Diff} = \left(D_{t+1} - D_t\,,\ D_{t+2} - D_{t+1}\,, \ldots, D_{t+(W-1)} - D_{t+(W-2)}\right)$, where W is the window size and D_t is the total electricity consumption for time t.

The design of this method also meant that the model would not recognise when the appliances were also being turned off, as highlighted in section 5.4.6.6. This is because the calculation for this window design gave positive values when an appliance was switched on and negative values when it was switched off. The feature set was calculated from these differences data, an example of this is shown in figure 5.17.

Electricity consumption data
221
213
219
214
217
215
212
214

Window number	Raw window data (reference data)						Difference data (reference data)					Average	Peak	Standard deviation	Root mean square
1	221	213	219	214	217	215	-8	6	-5	3	-2	-1.2	6	5.7184	5.2536
2	213	219	214	217	215	212	6	-5	3	-2	-3	-0.2	6	4.5497	4.0743
3	219	214	217	215	212	214	-5	3	-2	-3	2	-1	3	3.3912	3.1937

Figure 5.17: The transformation of the electricity data to the window data and feature set for trial 2– this figure shows three rows of consecutive sliding windows of window size 6.

For this analysis, three of the feature set values were excluded, these were the peak to average, peak to root mean square and root mean square to average ratios. These values were excluded from the analysis because they were calculated as infinity when either the average or the peak was 0. When the new feature set was calculated, the next steps of the analysis followed the same method as that used for the first trial as described in section 5.4.5. The recognition results from this analysis for the different window sizes with the four appliances that the model was trained to recognise is shown in table 5.7.

5.4.7.1. Trial 2- results

Window size	Appliance	True Positive	False Positive	False negative	True negative	Sensitivity	Specificity	Positive predictive value	Negative predictive value	Overall PPV	Overall Sensitivity
3	Shower	3	3	0	39280	100%	99.992%	50%	100%	2.88%	72.73%
	Microwave	3	94	1	39188	75%	99.761%	3.093%	99.997%		
	Kettle	2	169	0	39115	100%	99.570%	1.170%	100%		
	Dishwasher	0	4	2	39280	0%	99.990%	0%	99.995%		
4	Shower	3	6	0	39277	100%	99.985%	33.33%	100%	**3.79%**	**72.73%**
	Microwave	3	53	1	39229	75%	99.865%	5.357%	99.997%		
	Kettle	1	69	1	39215	50%	99.824%	1.429%	99.997%		
	Dishwasher	1	75	1	39209	50%	99.809%	1.316%	99.997%		
5	Shower	3	9	0	39274	100%	99.977%	25%	100%	2.75%	72.73%
	Microwave	3	47	1	39235	75%	99.880%	6%	99.997%		
	Kettle	1	129	1	39155	50%	99.672%	0.769%	99.997%		
	Dishwasher	1	98	1	39186	50%	99.751%	1.010%	99.997%		
6	Shower	3	12	0	39271	100%	99.969%	20%	100%	2.45%	72.73%
	Microwave	3	62	1	39220	75%	99.842%	4.615%	99.997%		
	Kettle	1	120	1	39164	50%	99.695%	0.826%	99.997%		
	Dishwasher	1	124	1	39160	50%	99.684%	0.8%	99.997%		
7	Shower	3	15	0	39268	100%	99.962%	16.667%	100%	2.69%	72.73%
	Microwave	3	56	1	39226	75%	99.857%	5.085%	99.997%		
	Kettle	1	68	1	39216	50%	99.827%	1.449%	99.997%		
	Dishwasher	1	150	1	39134	50%	99.618%	0.662%	99.997%		
8	Shower	3	18	0	39265	100%	99.954%	14.286%	100%	3.05%	81.82%
	Microwave	4	57	0	39225	100%	99.855%	6.557%	100%		
	Kettle	1	39	1	39245	50%	99.901%	2.5%	99.997%		
	Dishwasher	1	172	1	39112	50%	99.562%	0.578%	99.997%		
9	Shower	3	21	0	39262	100%	99.947%	12.5%	100%	1.27%	81.82%
	Microwave	3	367	1	38915	75%	99.066%	0.811%	99.997%		
	Kettle	1	36	1	39248	50%	99.908%	2.703%	99.997%		
	Dishwasher	2	275	0	39009	100%	99.300%	0.722%	100%		
10	Shower	3	24	0	39259	100%	99.939%	11.111%	100%	0.87%	72.73%
	Microwave	3	534	1	38748	75%	98.641%	0.559%	99.997%		
	Kettle	0	46	2	39238	0%	99.883%	0%	99.995%		
	Dishwasher	2	308	0	38976	100%	99.216%	0.645%	100%		

Table 5.7: The results from trial analysis 2 for window sizes 3 to 10

As can be seen from table A4.1 in appendix four, the window size that gave the best results for the first trial was window size 6, with an overall PPV of 1.8% and an overall sensitivity of 72.73%. For the results from trial 2 (shown in table 5.7) the window size that gave the best overall PPV and sensitivity was window size 4 with an overall PPV of 3.79% and an overall sensitivity of 72.73%.

Comparing the results from trial 2 (table 5.7) and the results from the first trial (shown in table A4.1, appendix four), the method used for trial 2 did give a similar or slightly higher overall sensitivity for all window sizes. This method does also slightly improve on the overall PPV for most window sizes when compared to the overall PPV scores from the first trial.

Analysing the results across all the window sizes, this method that was used for trial 2 did give a high overall sensitivity but a very low overall PPV percentage. The reasons why this method gave a very low PPV were examined further and the reason is discussed in the next sections.

5.4.7.2. Trial 2- discussion

The reason for the low overall PPV percentages is that this method gave a large number of false positives for all of the appliances. By analysing the results from window size 4, a problem was highlighted with the calculations for this window design method. The problem is highlighted by the results shown in table 5.8, this table shows the results from the model when it indicates that the shower was switched on.

Window number	Raw data (reference only)				Difference data (for reference only)			Average	Standard deviation	Peak	Root mean square
9664	139	240	219	9371	101	-21	9152	3077.333	5261.169	9152	5284.245
9665	240	219	9371	9376	-21	9152	5	3045.333	5288.544	9152	5283.924
9666	219	9371	9376	9327	9152	5	-49	3036	5296.680	9152	5283.986
10331	877	1184	818	9389	307	-366	8571	2837.333	4976.890	8571	4956.149
10332	1184	818	9389	9339	-366	8571	-50	2718.333	5071.020	8571	4953.063
10333	818	9389	9339	9293	8571	-50	-46	2825	4976.182	8571	4948.625
77215	347	349	371	9465	2	22	9094	3039.333	5243.505	9094	5250.439
77216	349	371	9465	9466	22	9094	1	3039	5243.794	9094	5250.439
77217	371	9465	9466	9421	9094	1	-45	3016.667	5263.175	9094	5250.488

Table 5.8: The results from the recognition model for the shower

The highlighted rows in table 5.8 show the true positives, i.e., when the model correctly predicted that the shower had been turned on. By comparing the results in the shaded rows with the results in the other (non-shaded) rows in the table, it can be seen that the model also classed the two immediately previous windows as the shower being turned on. This was because the value of the feature for all of the results (true positive and false positive) were all very similar due to the feature set being calculated from the difference data. This problem affected all of the appliances and explains the higher number of false positives for this method compared with the previous method.

To reduce the number of false positives, a post analysis filter was applied to the results from the recognition model. This ensured that only the results that had the power increase between the first and seconds points of the difference data set were included in the final recognition results. The method for this window design was re-run with this added filter and the results from the recognition model are shown below in table 5.9.

Window size	Appliance	True Positive	False Positive	False Negative	True Negative	Sensitivity	Specificity	Positive predictive value	Negative predictive value	Overall PPV	Overall Sensitivity
3	Shower	3	0	0	39283	100%	100%	100%	100%	5.56%	72.73%
	Microwave	3	48	1	39234	75%	99.878%	5.882%	99.997%		
	Kettle	2	86	0	39198	100%	99.781%	2.273%	100%		
	Dishwasher	0	2	2	39282	0%	99.995%	0%	99.995%		
4	Shower	3	0	0	39283	100%	100%	100%	100%	11.59%	72.73%
	Microwave	3	10	1	39272	75%	99.975%	23.077%	99.997%		
	Kettle	1	26	1	39258	50%	99.934%	3.704%	99.997%		
	Dishwasher	1	25	1	39259	50%	99.936%	3.846%	99.997%		
5	Shower	3	0	0	39283	100%	100%	100%	100%	9.09%	72.73%
	Microwave	3	7	1	39275	75%	99.982%	30%	99.997%		
	Kettle	1	48	1	39236	50%	99.878%	2.041%	99.997%		
	Dishwasher	1	25	1	39259	50%	99.936%	3.846%	99.997%		
6	Shower	3	0	0	39283	100%	100%	100%	100%	10.39%	72.73%
	Microwave	3	9	1	39273	75%	99.977%	25%	99.997%		
	Kettle	1	33	1	39251	50%	99.916%	2.941%	99.997%		
	Dishwasher	1	27	1	39257	50%	99.931%	3.571%	99.997%		
7	Shower	3	0	0	39283	100%	100%	100%	100%	15.09%	72.73%
	Microwave	3	8	1	39274	75%	99.980%	27.273%	99.997%		
	Kettle	1	12	1	39272	50%	99.969%	7.692%	99.997%		
	Dishwasher	1	25	1	39259	50%	99.936%	3.846%	99.997%		
8	Shower	3	0	0	39283	100%	100%	100%	100%	**20.93%**	**81.82%**
	Microwave	4	4	0	39278	100%	99.990%	50%	100%		
	Kettle	1	5	1	39279	50%	99.987%	16.667%	99.997%		
	Dishwasher	1	25	1	39259	50%	99.936%	3.846%	99.997%		
9	Shower	3	0	0	39283	100%	100%	100%	100%	8.11%	81.82%
	Microwave	3	65	1	39217	75%	99.835%	4.412%	99.997%		
	Kettle	1	4	1	39280	50%	99.990%	20%	99.997%		
	Dishwasher	2	33	0	39251	100%	99.916%	5.714%	100%		
10	Shower	3	0	0	39283	100%	100%	100%	100%	7.55%	72.73%
	Microwave	3	61	1	39221	75%	99.845%	4.688%	99.997%		
	Kettle	0	5	2	39279	0%	99.987%	0%	99.995%		
	Dishwasher	2	32	0	39252	100%	99.919%	5.882%	100%		

Table 5.9: The results from trial analysis 2 with filter, for window sizes 3 to 10

As can be seen from the results in table 5.9, the addition of the filter improved the overall PPV results across all windows compared with the results without the filter that are shown in table 5.7. The window size that gave the best overall PPV and sensitivity, with the filter, is window size 8 with an overall PPV of 20.93% and an overall sensitivity of 81.82%. This is a considerable improvement in the overall PPV when compared to the best results without the filter, which had an overall PPV of 3.79%. The results for the overall sensitivity also improved from 72.73% to 81.82%.

Although applying a filter to the results did decrease the number of false positives and increase the values of the overall PPV and sensitivity, this method still produced too many false positives and therefore the overall PPV was not acceptable. To try and decrease the numbers of false positives, a further method was designed.

5.4.8. Trial 3

To address the problem with the number of false positives that were produced by the method used in trial 2 a new window design was developed that incorporated the two previous window designs. The average, peak, standard deviation and root mean square values were calculated from the difference data as in trial 2. The peak to average ratio, peak to root mean square ratio and the root mean square to average ration were calculated from the window data as in the first trial. This new window design is shown in figure 5.18.

Electricity consumption data
221
213
219
214
217
215
212
214

Window number	Raw window data (reference only)						Difference data (reference only)					Average	Peak	Standard deviation	Root mean square	Peak to average ratio	Peak to root mean square ratio	Root mean square to average ratio
1	221	213	219	214	217	215	-8	6	-5	3	-2	-1.2	6	5.718	5.254	1.021	1.021	1.0001
2	213	219	214	217	215	212	6	-5	3	-2	-3	-0.2	6	4.550	4.074	1.019	1.019	1.0001
3	219	214	217	215	212	214	-5	3	-2	-3	2	-1.0	3	3.391	3.194	1.018	1.018	1.0001

Figure 5.18: The transformation of the electricity data to the window data and feature set for trial 3– this figure shows three rows of consecutive sliding windows of window size 6.

Now that the new feature set had been calculated, the next steps of the analysis followed the same method as that used for trial 2, with the addition of the post analysis filter, as described in section 5.4.7. The recognition results from this analysis for the different window sizes with the four appliances that the model was trained to recognise is shown in table 5.10. For reference, this method was also run without the addition of a post analysis filter, the results from this are in table A4.2, appendix four.

5.4.8.1. Trial 3- results and discussion

Window size	Appliance	True Positive	False Positive	False negative	True negative	Sensitivity	Specificity	Positive predictive value	Negative predictive value	Overall PPV	Overall Sensitivity
3	Shower	3	0	0	39283	100.0%	100%	100%	100%	2.70%	72.73%
	Microwave	3	83	1	39199	75.0%	99.789%	3.488%	99.997%		
	Kettle	2	203	0	39081	100.0%	99.483%	0.976%	100%		
	Dishwasher	0	2	2	39282	0.0%	99.995%	0.000%	99.995%		
4	Shower	3	0	0	39283	100.0%	100%	100%	100%	4.35%	72.73%
	Microwave	3	4	1	39278	75.0%	99.990%	42.857%	99.997%		
	Kettle	1	146	1	39138	50.0%	99.628%	0.680%	99.997%		
	Dishwasher	1	26	1	39258	50.0%	99.934%	3.704%	99.997%		
5	Shower	3	0	0	39283	100.0%	100%	100%	100%	5.37%	72.73%
	Microwave	3	3	1	39279	75.0%	99.992%	50.000%	99.997%		
	Kettle	1	111	1	39173	50.0%	99.717%	0.893%	99.997%		
	Dishwasher	1	27	1	39257	50.0%	99.931%	3.571%	99.997%		
6	Shower	3	0	0	39283	100.0%	100%	100%	100%	10.53%	72.73%
	Microwave	3	3	1	39279	75.0%	99.992%	50%	99.997%		
	Kettle	1	38	1	39246	50.0%	99.903%	2.564%	99.997%		
	Dishwasher	1	27	1	39257	50.0%	99.931%	3.571%	99.997%		
7	Shower	3	0	0	39283	100.0%	100%	100%	100%	18.60%	72.73%
	Microwave	3	1	1	39281	75.0%	99.997%	75%	99.997%		
	Kettle	1	8	1	39276	50.0%	99.980%	11.111%	99.997%		
	Dishwasher	1	26	1	39258	50.0%	99.934%	3.704%	99.997%		
8	Shower	3	0	0	39283	100.0%	100%	100%	100%	19.51%	72.73%
	Microwave	3	3	1	39279	75.0%	99.992%	50%	99.997%		
	Kettle	1	5	1	39279	50.0%	99.987%	16.667%	99.997%		
	Dishwasher	1	25	1	39259	50.0%	99.936%	3.846%	99.997%		
9	Shower	3	0	0	39283	100.0%	100%	100%	100%	**20.41%**	**90.91%**
	Microwave	4	6	0	39276	100.0%	99.985%	40%	100%		
	Kettle	1	10	1	39274	50.0%	99.975%	9.091%	99.997%		
	Dishwasher	2	23	0	39261	100.0%	99.941%	8%	100%		
10	Shower	3	0	0	39283	100.0%	100%	100%	100%	15.69%	72.73%
	Microwave	3	7	1	39275	75.0%	99.982%	30%	99.997%		
	Kettle	0	9	2	39275	0.0%	99.977%	0%	99.995%		
	Dishwasher	2	27	0	39257	100.0%	99.931%	6.897%	100%		

Table 5.10: The results from trial analysis 3 for window sizes 3 to 10

Comparing the results in tables 5.10 and (table A4.2, appendix four), i.e., the results for the same methods but without and with the post analysis filter it is clear that the addition of the post analysis filter for this method again improved the overall PPV results across all the windows. The best overall PPV was for a window size of 9, with an overall PPV of 20.41% compared to a best overall PPV of 5.97% for the results without a post analysis filter. The best overall sensitivity also improved with the addition of a filter from 72.73% (table A4.2, appendix four) to 90.91%.

Comparing the results in table 5.10 with the results from the previous trial that are shown in table 5.9, this method produced a slightly worse result in terms of overall PPV. With the best overall PPV given by window size 9 with an overall PPV 20.43% compared to 20.93% for trial 2. However, this trial did give a higher overall sensitivity of 90.91%, compared with 81.82% for trial 2. This method still produced a large number of false positives and did not improve on the previous trial results. To try and decrease the numbers of false positives produced a further method was designed, and tested in Trial 4.

5.4.9. Trial 4

To try and improve on the results from the previous method a new window was designed that followed the same method as the previous window but all the data were calculated using the difference data, rather than the average, peak, standard deviation and root mean square being calculated using the difference data and the peak to average ratio, peak to root mean square ratio and the root mean square to average ratio being calculated from the normal window data (as described in trial 1).

Electricity consumption data
221
213
219
214
217
215
212
214

Window number	Raw window data (reference only)	Difference data (reference only)	Average	Peak	Standard deviation	Root mean square	Peak to average ratio	Peak to root mean square ratio	Root mean square to average ratio
1	221 213 219 214 217 215	-8 6 -5 3 -2	-1.2	6	5.718	5.254	-5	1.142	-4.378
2	213 219 214 217 215 212	6 -5 3 -2 -3	-0.2	6	4.550	4.074	-30	1.473	-20.372
3	219 214 217 215 212 214	-5 3 -2 -3 2	-1.0	3	3.391	3.194	-3	0.939	-3.194

Figure 5.19: The transformation of the electricity data to the window data and feature set for trial 4– this figure shows three rows of consecutive sliding windows of window size 6

For some of the window rows, the average or the peaks values equalled 0. These rows were excluded from the data set that was used to train the model as the peak to average, peak to root mean square or the root mean square to average values had led to values of infinity for those data points.

Following this change, the new feature set was calculated and this meant that for all the rows that had previously contained infinity values, these were no longer present. The next steps of the analysis follows the same method as that used for the trial 2 with the addition of a post analysis filter, as described in section 5.4.7. The recognition results from this analysis for the different window sizes with the four appliances that the model was trained to recognise is shown in table 5.11. For reference, this method was also run without the addition of a post analysis filter, the results from this are in table A4.3, appendix four.

5.4.9.1. Trial 4- results and discussion

Window size	Appliance	True Positive	False Positive	False Negative	True Negative	Sensitivity	Specificity	Positive predictive value	Negative predictive value	Overall PPV	Overall Sensitivity
3	Shower	3	0	0	31387	100%	100%	100%	100%	7.69%	72.73%
	Microwave	3	57	1	31329	75%	99.818%	5.0%	99.997%		
	Kettle	2	37	0	31351	100%	99.882%	5.128%	100%		
	Dishwasher	0	2	2	31386	0%	99.994%	0%	99.994%		
4	Shower	3	0	0	31797	100%	100%	100%	100%	14.81%	72.73%
	Microwave	3	9	1	31787	75%	99.972%	25.000%	99.997%		
	Kettle	2	17	0	31781	100%	99.947%	10.526%	100%		
	Dishwasher	0	20	2	31778	0%	99.937%	0%	99.994%		
5	Shower	3	0	0	32346	100%	100%	100%	100%	11.94%	72.73%
	Microwave	3	8	1	32337	75%	99.975%	27.273%	99.997%		
	Kettle	2	41	0	32306	100%	99.873%	4.651%	100%		
	Dishwasher	0	10	2	32337	0%	99.969%	0.0%	99.994%		
6	Shower	3	0	0	32683	100%	100%	100%	100%	**18.75%**	**81.82%**
	Microwave	3	5	1	32677	75%	99.985%	37.50%	99.997%		
	Kettle	2	15	0	32669	100%	99.954%	11.765%	100%		
	Dishwasher	1	19	1	32665	50%	99.942%	5%	99.997%		
7	Shower	3	0	0	33074	100%	100%	100%	100%	18.18%	72.73%
	Microwave	3	6	1	33067	75%	99.982%	33.333%	99.997%		
	Kettle	2	16	0	33059	100%	99.952%	11.111%	100%		
	Dishwasher	0	14	2	33061	0%	99.958%	0%	99.994%		
8	Shower	3	0	0	33265	100%	100%	100%	100%	19.51%	72.73%
	Microwave	3	4	1	33260	75%	99.988%	42.857%	99.997%		
	Kettle	2	24	0	33242	100%	99.928%	7.692%	100%		
	Dishwasher	0	5	2	33261	0%	99.985%	0%	99.994%		
9	Shower	3	0	0	33672	100%	100%	100%	100%	10.34%	54.55%
	Microwave	0	21	4	33650	0%	99.938%	0%	99.988%		
	Kettle	1	8	1	33665	50%	99.976%	11.111%	99.997%		
	Dishwasher	2	23	0	33650	100%	99.932%	8%	100%		
10	Shower	3	0	0	33940	100%	100%	100%	100%	9.72%	63.64%
	Microwave	0	39	4	33900	0%	99.885%	0%	99.988%		
	Kettle	1	8	1	33933	50%	99.976%	11.111%	99.997%		
	Dishwasher	2	19	0	33922	100%	99.944%	9.524%	100%		

Table 5.11: The results from trial analysis 4 for window sizes 3 to 10

As shown by the results in table 5.11, the addition of the filter has improved the overall PPV results across all windows compared with the results without the filter that are shown in table A4.3, appendix four. For the results with the filter, the best overall PPV and sensitivity were given by window size 8 but this window did not detect any true positives for the dishwasher. Therefore the best window size that gave true positives for all the appliances is window size 6, with an overall PPV of 18.75% and an overall sensitivity of 81.82% This was a considerable improvement in the overall PPV from the results without the post analysis filter for which the best overall PPV was 3.66%, with an overall sensitivity of 81.82%.

Comparing the results in table 5.11 with the results from the previous two trials (tables 5.9 and 5.10), this method gave the worst results, with the best overall PPV for trial 4 being 18.75% compared with 20.41% from trial 3 and 20.93% from trial 2. For the overall sensitivity this trial gave a best overall sensitivity of 81.82% this was the same as trial 2 but worse than the best overall sensitivity 90.91% given by trial 3.

5.4.10. Conclusion

Table 5.12 provides an overview of the 4 different types of window design that were used within this section of results. This table also gives the best window size and overall PPV and sensitivity for each of the different trial window designs.

Trial	Feature set	Best window size	Best overall PPV	Best overall sensitivity
1	Average Peak Standard deviation Root mean square Peak to average ratio Peak to root mean square ratio Root mean square to average ratio	6	1.80%	72.73%
2	Difference average Difference peak Difference standard deviation Difference root mean square	4	3.79%	72.73%
2 with filter	Difference average Difference peak Difference standard deviation Difference root mean square	8	20.93%	81.82%
3	Difference average Difference peak Difference standard deviation Difference root mean square Peak to average ratio Peak to root mean square ratio Root mean square to average ratio	8	5.97%	72.73%
3 with filter	Difference average Difference peak Difference standard deviation Difference root mean square Peak to average ratio Peak to root mean square ratio Root mean square to average ratio	9	20.41%	90.91%
4	Difference average Difference peak Difference standard deviation Difference root mean square Difference peak to average ratio Difference peak to root mean square ratio Difference root mean square to average ratio	6	3.66%	81.82%
4 with filter	Difference average Difference peak Difference standard deviation Difference root mean square Difference peak to average ratio Difference peak to root mean square ratio Difference root mean square to average ratio	6	18.75%	81.82%

Table 5.12: Overview of results from the 4 trials

As the results in table 5.12 show, the three window designs with the post analysis filter, as described in section 5.4.7, 5.4.8, 5.4.9, gave the best overall PPV and best overall sensitivity. For these results the window size that gave the best overall PPV and best overall sensitivity for these results varied, i.e., the best results being given by window size 8 for trial 2, window size 9 for trial 3 and window size 6 for trial 4. Therefore, it was decided that the results for the next stages of the analysis would be undertaken again using multiple window sizes.

The results from the three window designs which all used the post analysis filter, all gave approximately the same overall PPV of around 20% and varying overall sensitivities of between 81.82% and 90.91%. Although the values for the overall sensitivities for these window designs were acceptable, the overall PPV values were not acceptable and a new window design needed to be developed.

5.5. Appliance recognition- window re-design

As described in section 5.4, the results given by the different trials did not produce an overall PPV of greater than 20.93%; therefore, a new method was developed. This section will give a description of the design of this method and the results obtained.

5.5.1. Window re-design

As described in section 5.4.1, the previous design of the window and feature sets used in section 5.4 were based on the method used by (Lee et al., 2010; Lin et al., 2010). For this method of window design and the window used so far for this project, the windows only looked forward in the data. It was decided to redesign the window design from the work previously done so as to develop a window design that also looked back a certain number of instances as well as forwards. Figure 5.20 gives an overview of how this backwards and forward-looking window was designed.

Electricity consumption data
221
213
219
214
217
215
212
214
213

Window

Window number	Raw window data							Difference data (2 back-4 forward)					
1	221	213	219	214	217	215	212	(221-219)	(213-219)	(214-219)	(217-219)	(215-219)	(212-219)
2	213	219	214	217	215	212	214	(213-214)	(219-214)	(217-214)	(215-214)	(212-214)	(214-214)
3	219	214	217	215	212	214	213	(219-217)	(214-217)	(215-217)	(212-217)	(214-217)	(213-217)

Feature set

Window number	Raw window data (reference only)							Difference data for 2 back – 4 forward (reference only)						Average	Peak	Standard deviation	Root mean square	Peak to average ratio	Peak to root mean square ratio	Root mean square to average ratio
1	221	213	219	214	217	215	212	2	-6	-5	-2	-4	-7	-3.7	2	3.266	4.726	-0.55	0.423	-1.289
2	213	219	214	217	215	212	214	-1	5	3	1	-2	0	1.0	5	2.608	2.582	5	1.936	2.582
3	219	214	217	215	212	214	213	2	-3	-2	-5	-3	-4	-2.5	2	2.429	3.342	-0.8	0.599	-1.337

Figure 5.20: The transformation of the electricity data to the window data and feature set – this figure shows three rows of consecutive sliding windows of 2 backwards and 4 forwards.

As shown in figure 5.20 there are three steps needed to create a backwards and forwards window design. This first step of this design was to create the windows, which contained 2 backward points and 4 forward points. In the example shown in the raw window data column (shown in the second and third tables in figure 5.20), there were actually 7 data points in the raw window data. The seventh data point, or the reference point, is the data point that gives the window design, for this example, of 2 backwards and 4 forward from its point in the window. The value of this reference point was then subtracted from all of the data points of raw window data (as shown by the difference data column of the second table in figure 5.20). The reference point column was then removed from the difference data set as it would give a value of zero and the final results from this calculation are shown in the difference data in column in the third table in figure 5.20. It was the values of these results that were then used to calculate the values of the feature set, as shown by the third table in figure 5.20. The feature set used for this window redesign is the same as that used for all the previous window design.

5.5.2. Trial 5

To try and improve on the results from the previous trials, trial 5 was run using the feature set calculated using the new window design as described in section 5.5.1. As well as the changes to the calculation of the feature set, this trial also investigated the combinations of variables in the feature set that produced the best results, in terms of overall PPV and sensitivity. To achieve this, a loop was created that ran different combinations of the feature set variables, for each of the window sizes from back 1-6 and forwards 2-6. The list of the different feature set combinations that were run is shown in table A5.1 of appendix five. Running the data for window size back 1-6 and forwards 2-6 and for the different combinations, produced 2970 sets of results. This is too much detail to incorporate into the thesis, so only the most promising results from this trial will be shown. The full results from this trial are available on request from the author.

For this trial all of the appliances that this household was able to record were considered. There were four appliances that were excluded from this trial. The four appliances were the kettle, toaster, television and electric hob with the reasons for their exclusions discussed in the sections below.

5.5.2.1. Trial 5- the kettle

The change in the way the window, and therefore the feature set, was calculated highlighted a problem with both the data used to train the kettle but also with how the training and test data were collected.

The aim of this project was to produce a non-intrusive method for monitoring resident’s activities via their electricity usage and so to fit in with the aim of being non-intrusive the residents who took part in this data collection were not asked to change from their daily routines. Therefore, because the test and training data for each appliance were not collected in isolation there was a risk for some appliance signatures being 'distorted' by other appliances turning on or being on at the same time. This problem was particularly evident with the training and test data collected from the kettle, where it was almost always turned on with or when other appliances were on.

Although the aim of this forward and backwards window design was to limit the disruption caused by other appliances being on in the background, for the case of the kettle, this window design highlighted too many discrepancies with the training data set. Figure 5.21 to figure 5.23 show 3 instances of when the kettle is turned on (based on the diary data) and the corresponding electricity consumption data for that time.

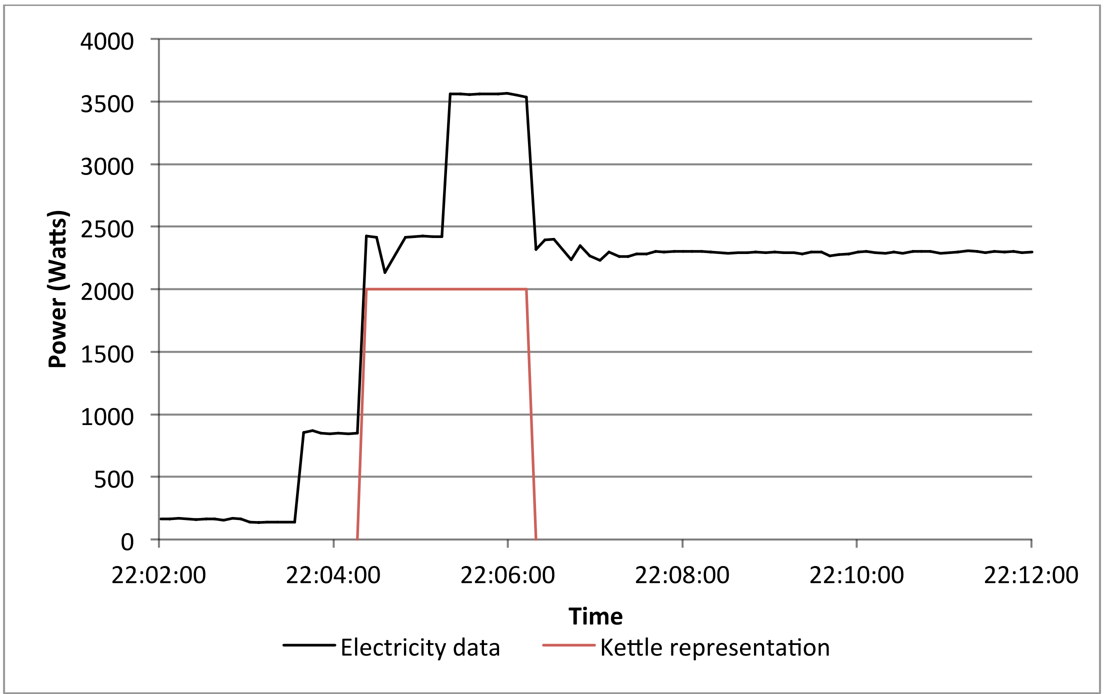

Figure 5.21: Electricity consumption data with a representation of the kettle

From the electricity data in figure 5.21, another appliance came on just before the kettle as shown by the power rise from about 200 Watts to about 800 Watts. This distorted the kettle signature and therefore the feature set that was used to train the model.

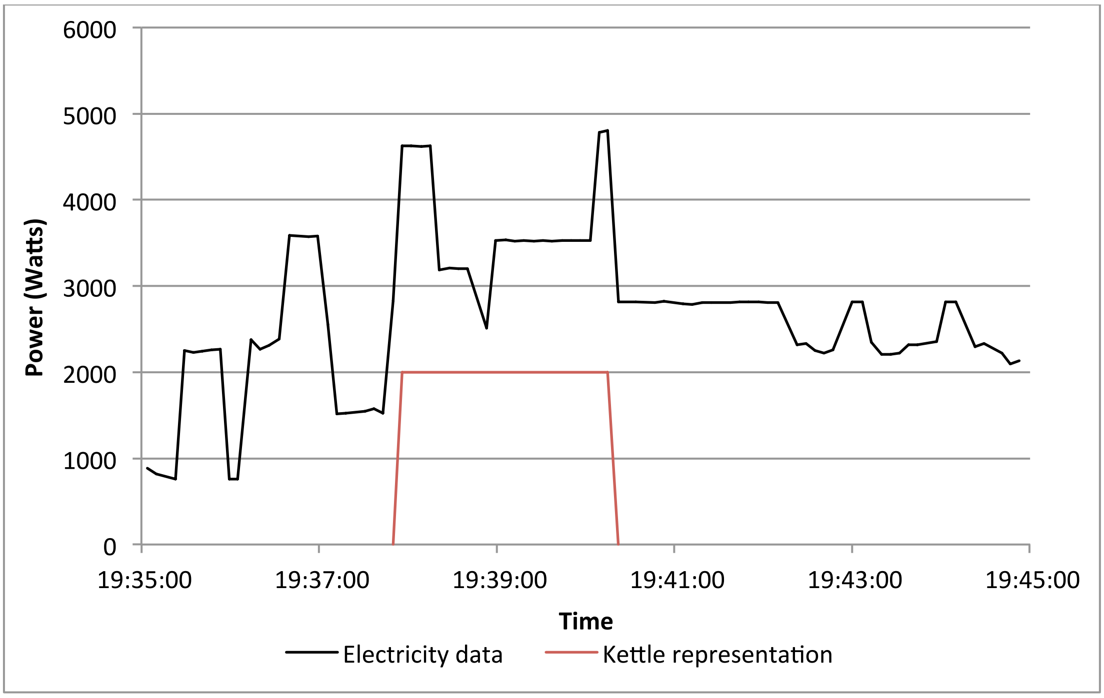

Figure 5.22: Electricity consumption data with a representation of the kettle

For the electricity data shown in figure 5.22, there is already an appliance on when the kettle is turned on. As the appliance and the kettle are turned on at almost the same time their signatures combine. This is represented by the power rise of about 1500 Watts to 4500 Watts. As the power rating of the kettle for this house is about 2000 Watts, this does not represent the kettle for this house very well.

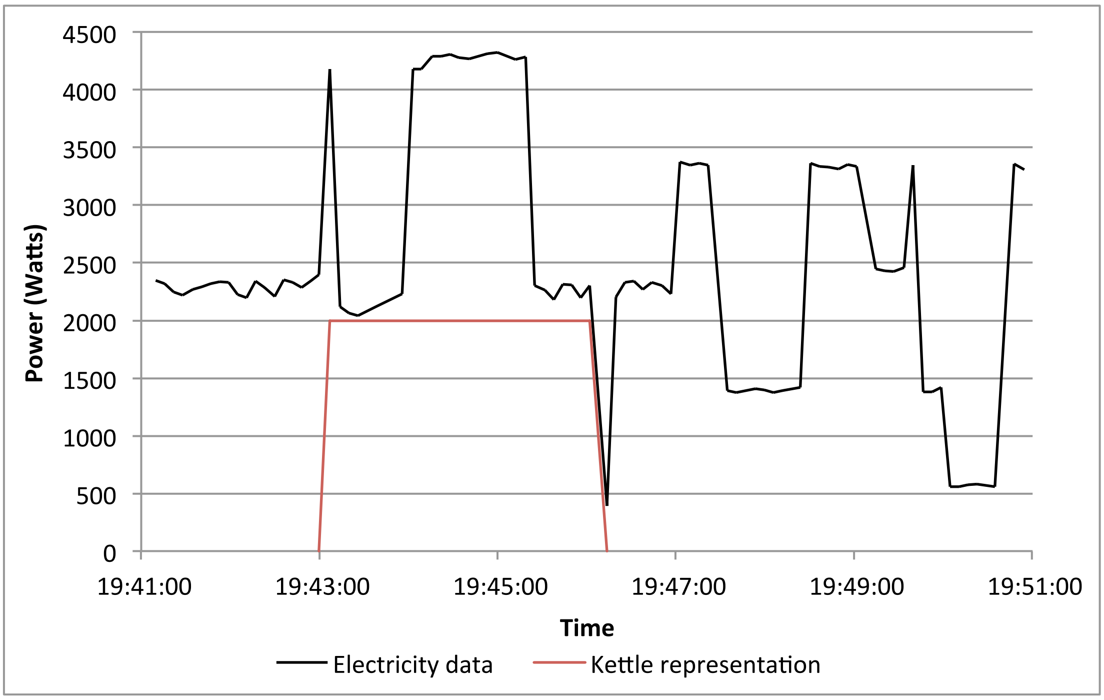

Figure 5.23: Electricity consumption data with a representation of the kettle

The electricity data shown in figure 5.23 does not seem to represent that a kettle has been turned on apart from a spike rise from about 2200 Watts to about 4200 Watts. This could be explained as another appliance of a similar power turning off just after the kettle has been turned on and masking the kettle's signature.

With the way that this window design "looks" both forward and backwards, if there are large changes in power around the point when an appliance is turned on, this can undermine the feature set. All of the three instances shown in figures 5.21-5.23 have other appliances on or turned on at approximately the same time; therefore, the other appliances distort the feature set associated with the kettle, which would be used to train the model. The aim when selecting training data for modelling is to provide a feature set that gives a good representation of the features associated with the "object" that one wants to train the model to recognise.

In the case of the kettle, the data collected had too many discrepancies within its training feature set, for the reason discussed above, and did not provide a good representation of the features associated with the kettle. It was therefore decided to remove the kettle from the recognition model for this house, as the data were not sufficient to provide an accurate and robust recognition model.

5.5.2.2. Trial 5- the toaster

Although this is a large amount of data, when translated down to appliance usage, in the case of the toaster it only meant two instances that the toaster has been used. This meant that it made sense to exclude the toaster, which was recorded by this household, due to the limited amount of data for training and testing.

5.5.2.3. Trial 5- the electric hob

For this household the instances of use for the electric hob were recorded and there were sufficient instances of use to provide test and training data. When recording the hob data, the participants were only asked to record when they had used the appliance not the number of hobs used or the power of each of these. However, as hobs have multiple sections and multiple powers, the electrical signatures for the hobs can be variable. This variability was highlighted in the test and training data for the hob, where the feature set values for each instance were noticeably different from each other. The conclusion from this is that for each of the instances of the hob, there were different numbers of hobs and or different power ratings used. With this variability with the number of different possibilities for the signatures of the hob and the corresponding variability in the values of the feature set the hob was excluded from this trial. This was due to there being insufficient data collected to be able to train an accurate and robust recognition model that could recognise all the different instances of hob use.

5.5.2.4. Trial 5- the television

The television for this household was also recorded and there were enough instances of use to provide training and test data sets. However, from the training and test data sets, the power rating for the television was found to be very low (<40Watts). This low power made it impossible to distinguish the television above background the noise of the data. For this reason, the television was excluded from this trial due to an insufficiently high ‘signal’.

5.5.3. Trial 5- results and discussion

From analysing the results given by all the window sizes and feature set combinations, the window size and combination that gave the best results in terms of overall PPV was backwards 6 and forwards 2, using the features of average, peak to average ratio and root mean square to average ratio. For this window size and feature set the overall PPV was 55.55%, although for this window size and combination, the overall sensitivity was 38.46%. As this sensitivity was not satisfactory for the aims of this method, the results were analysed again to find the best overall PPV with a satisfactory overall sensitivity.

From this analysis, there were eight iterations that gave a result with a high overall PPV and a high overall sensitivity. From these results there were multiple window sizes and feature sets that gave the same result. The results shown in the tables below are given in terms of true positives (TP), false positives (FP), false negatives (FN), true negatives (TN), sensitivity, specificity, positives predicted value (PPV) and negative predicted value (NPV) for each appliance. The overall PPV and overall sensitivity from all the appliances is also shown, along with the window size and feature set that gave the results.

Appliance	TP	FP	FN	TN	Sensitivity	Specificity	PPV	NPV	Overall PPV	Overall Sensitivity
Shower	3	0	0	38512	100%	100%	100%	100%	**28.57%**	**76.92%**
Microwave	3	1	1	38510	75%	99.997%	75%	99.997%		
Dishwasher	1	15	1	38498	50%	99.961%	6.25%	99.997%		
Oven	2	8	0	38505	100%	99.979%	20%	100%		
Washing machine	1	1	1	38512	50%	99.997%	50%	99.997%		

Table 5.13: Results from window size 6 backwards 4 forwards, feature set (standard deviation, root mean square, peak to average ratio)

Appliance	TP	FP	FN	TN	Sensitivity	Specificity	PPV	NPV	Overall PPV	Overall Sensitivity
Shower	3	0	0	38414	100%	100%	100%	100%	**28.57%**	**76.92%**
Microwave	3	1	1	38412	75%	99.997%	75%	99.997%		
Dishwasher	1	16	1	38399	50%	99.958%	5.882%	99.997%		
Oven	2	7	0	38408	100%	99.982%	22.222%	100%		
Washing Machine	1	1	1	38414	50%	99.997%	50%	99.997%		

Table 5.14: Results from window size 6 backwards 5 forwards

For the results in table 5.14 there were five feature sets that gave the same results with the same window size 6 backwards 5 forwards for each. The first feature set was average, peak, standard deviation and peak to average ratio. The second feature set was average, peak, root mean square and peak to average ratio. The third feature set was average, standard deviation root mean square and peak to average ratio. The fourth feature set was average, peak, standard deviation, root mean square, and peak to average ratio. The fifth feature set was average, standard deviation, root mean square, peak to average ratio and peak to root mean square ratio.

Appliance	TP	FP	FN	TN	Sensitivity	Specificity	PPV	NPV	Overall PPV	Overall Sensitivity
Shower	3	0	0	38512	100%	100%	100%	100%	**28.57%**	**76.92%**
Microwave	3	1	1	38510	75%	99.997%	75%	99.997%		
Dishwasher	1	17	1	38498	50%	99.956%	5.556%	99.997%		
Oven	2	6	0	38505	100%	99.984%	25%	100%		
Washing machine	1	1	1	38512	50%	99.997%	50%	99.997%		

Table 5.15: Results from window size 6 backwards 5 forwards

For the results in table 5.15 there were two feature sets that gave the same results with the same window size 6 backwards 5 forwards for each. The first feature set was peak, standard deviation, root mean square and peak to average ratio. The second feature set was peak, standard deviation, root mean square, peak to average ratio and peak to root mean square ratio.

As shown by the results in tables 5.13, 5.14 and 5.15, the best results with a high overall PPV and high overall sensitivity were an overall PPV of 28.57% and an overall sensitivity of 76.92%. Comparing these results with those from all of the previous trials, the value of the overall PPV is a slight improvement from the best results of 20.93%, from trials 1-4 in sections 5.5.7-5.5.9.

The reason for the low value of overall PPV was due to the larger number of false positives, most of which were produced by the dishwasher and the oven. As the results from this trial were not acceptable due to the low value of the overall PPV, the results were analysed further to discover why the oven and dishwasher produced a large number of false positives. The reasons for this are discussed in the next section.

5.5.3.1. Trial 5- further analysis

From the results shown in tables 5.13, 5.14 and 5.15, the oven and the dishwasher produced a large number of false positives, which affected the results for the overall PPV. The reasons for both of these appliances producing a large number of false positives are discussed in more detail below, starting with the oven.

5.5.3.1.1. Trial 5- oven

Analysing why the oven produced a large number of false positives found that this model was classifying the oven cycle of heating up until the correct temperature was reached and the oven thermostat turning the heat off, as the oven turning on. The example signature of the oven, shown in figure 5.5, highlights why this model may identify each time when the oven's temperature has fallen so the heating element comes on as a new event of the oven being turned on. As figure 5.5 shows, the oven repeats are of an almost exactly the same power as when the oven was first turned on but of a different time period. As this method did not look at the time that appliance signatures were on for, the model recognised the repeats as if the appliance has just been turned on again. This therefore creates an "artificial" false positive in the results as the oven was being switched on again but by the thermostat of the oven and not the occupant.

To reduce the number of repeats of the oven that this model recognised, additional code was written that would ignore subsequent oven points if the model recognised the oven and there was a previous oven point within 15 minutes. For example as each of the oven "on" points and oven "repeats", shown in figure 5.5, were within 15 minutes of the next oven point only the first point, i.e., when the oven was first turned on, would be identified as oven the being turned on by this model. The reason why the time limit of 15 minutes was chosen was to cover all the instances, within an acceptable margin of error, of the oven first warming up and then turning on again for repeats. The reason for this margin of error was also so as to cover instances of the oven being turned up to higher temperatures, as the residents did not record the temperature the oven was set at.

5.5.3.1.2. Trial 5- dishwasher

As described in section 5.5.3.1.1, the dishwasher, similar to the oven, produced repeats that were of a similar size to when they were first turned on, as highlighted by the example signature in figure 5.8. To reduce the repeats of the dishwasher, similar code for that used for the oven repeats (section 5.5.3.1.1) was written. For the dishwasher, if the model recognised the dishwasher and there was a previous dishwasher point within 20 minutes of this dishwasher point, then this point would be ignored.

Another reason for the high number of false positives from the dishwasher was that the model recognised some of the instances of the oven "repeats" as the dishwasher. The reason for this is the similarity in signature and feature set of the oven "repeats" and dishwasher in this household.

5.5.3.1.3. Trial 5- washing machine

The washing machine was another appliance, similar to the oven and dishwasher, which produced repeats that were similar in size to when the appliance was first turned on, as highlighted in figure 5.7. However, for the best result from this trial, shown in tables 5.13 – 5.15, the repeats of the washing machine did not affect the overall result. Because the model only gave one false positive, there were other window sizes and feature set combinations that were affected by the repeats. To reduce the repeats of the washing machine, similar code for that used for the oven repeats (section 5.5.3.1.1) was written. For the washing machine, if the model recognised the washing machine and there was a previous washing machine point within 10 minutes of this washing machine point, then this point would be ignored.

5.5.4. Trial 6

This trial follows the same method as used for trial 5 but with the additions of the filters on the repeats of appliances that are described in section 5.5.3.1.1. The aim of this trial was to compare the results with trial 5, to see whether the addition of these filters would remove the repeats of appliances and improve the overall PPV results. The results from this trial are discussed in the section below.

5.5.5. Trial 6- results and discussion

As in the previous trial, this window design was run for multiple window sizes (back 1-6 and forwards 2-6) as well as different combinations of the window design's feature set. The list of the different feature set combinations that were run is shown in table A5.1 of appendix five: running the data for window size back 1-6 and forwards 2-6, and for the different combinations of these, produced 2970 sets of results. This is too many pages, and too much detail, to incorporate into the thesis, so only the most promising results from this trial will be shown. The full results from this trial are available on request from the author.

Analysing the results in terms of overall PPV, the best result of 71.43% was giving by window size, 6 back 4 forwards, using the features of average, peak to average ratio and peak to root mean square to average ratio. The overall sensitivity for this result was 38.46% although, for this result, the model did not recognise any instances of the dishwasher or the washing machine. This result was not seen as acceptable, as no instances of the oven, dishwasher or the washing machine were recorded, so the results were not analysed further.

There were two iterations that gave a result with a high overall PPV and a high overall sensitivity. The results for these are shown in tables 5.16, as both of the iterations gave exactly the same results, in terms of true positives, false positives etc. for each appliance. Both of the iterations had a window size of 6 backwards and 4 forwards, with the first feature set being peak, standard deviation and peak to average ratio. The second feature set was standard deviation, root mean square and peak to average ratio.

Appliance	TP	FP	FN	TN	Sensitivity	Specificity	PPV	NPV	Overall PPV	Overall Sensitivity
Shower	3	0	0	38414	100%	100%	100%	100%	**62.50%**	**76.92%**
Microwave	3	1	1	38412	75%	99.997%	75%	99.997%		
Dishwasher	1	2	1	38413	50%	99.995%	33.33%	99.997%		
Oven	2	2	0	38413	100%	99.995%	50%	100%		
Washing Machine	1	1	1	38414	50%	99.997%	50%	99.997%		

Table 5.16: Results from window size 6 backwards 4 forwards, feature set (peak, standard deviation and peak to average ratio) and feature set (was standard deviation, root mean square and peak to average ratio)

As shown by the results in tables 5.16, the best results, i.e., with a high overall PPV and high overall sensitivity, was an overall PPV of 62.50% and an overall sensitivity of 76.92%. Comparing these results with the results from the previous trial (trial 5), the removal of the repeats of the oven, dishwasher and washing machine improved the overall PPV results, with the best overall PPV for the previous trial (trial 5) being 28.57% compared with 62.50% for this trial. The overall sensitivity was the same for both runs of the trial.

The results from this trial were also an improvement on the best overall PPV results from trials 1-4, sections 5.5.7-5.5.9, which had a best overall PPV of 20.93%. For the overall sensitivity the result decreased from the best results from trials 1-4 of 90.91%.

5.5.5.1. Trial 6- random baseline results

The random baseline results from the best feature set and window size combination, as shown in table 5.16, for trial 6 are also provided for further comparisons of the results from this trial. The random baseline method used for these results followed the same method as described in section 5.4.5.1, with the results shown in table 5.17.

Appliance	TP	FP	FN	TN	Sensitivity	Specificity	PPV	NPV	Overall PPV	Overall Sensitivity
Shower	0	33	3	38381	0%	99.914%	0%	99.992%	**0%**	**0%**
Microwave	0	4	4	38409	0%	99.990%	0%	99.990%		
Dishwasher	0	3	2	38412	0%	99.992%	0%	99.995%		
Oven	0	4	2	38411	0%	99.990%	0%	99.995%		
Washing machine	0	0	2	38415	0%	100%	NaN	99.995%		

Table 5.17: Random baseline of best results from trial 6 with window size 6 backwards 4 forwards, feature set (peak, standard deviation and peak to average ratio)

These results provide a comparison between a model trained on randomly selected appliances points (random baseline) and the final trial model (section 5.5.5). As shown in table 5.17, the random baseline gave an overall PPV and sensitivity of 0%, this is a very poor performance when compared with the an overall PPV and sensitivity of 62.50% and 79.62% as provided from the best window and feature set combination in table 5.16.

For reference, the attribute values for each class are provided in appendix six, table A6.1 for the best window size and feature set combination (as shown in table 5.16). As shown in figure 5.20, the attributes are calculated for the window configuration: because this changes during each iteration, it would be too much information to show each attribute value for each class for each iteration.

5.5.6. Trial 6- cross validation

As discussed in section 5.4.4, the results shown in section 5.5.5 are from the use of the holdout method for dividing the data into a training and test datasets. As previously discussed in section 2.4.5.9, the use of this method has some disadvantages such as potential overfitting or not providing generalizable results, as the training or test data might not be representative of the wider data. A method for overcoming these concerns is to conduct cross validation; section 2.4.5.8 has highlighted and discussed some of the different cross-validation methods.

For this research, k-fold cross validation was chosen instead of leave-one-out cross-validation, because the latter method is computationally too expensive, due to the size of the data set. Due to the low usage of some of the appliances, the value of k was chosen to be three, because any larger number would not leave enough appliances points in each of the folds. As highlighted by Witten et al., (2011), previous research has suggested that for optimal results 10-fold cross validation should be used. This was not possible for this data due to the low level of appliance usage in the households. For the cross validation, as mentioned above, three-fold cross-validation was chosen; therefore, the data set was divided into three, randomly chosen, approximately equal-sized folds, with each fold having at least one data point from each of the appliances.

As highlighted by Witten et al., (2011), due to the variation in how the folds are chosen, different results from the cross validation will be produced. To provide a reliable estimate of the performance, cross validation can be carried out a number of times with the results averaged to produce the final result. For this trial, the three-fold cross validation was run three times for each window size and feature set, as undertaken in section 5.5.4, with the aim, as highlighted by Witten et al., (2011), of achieving a reliable estimate of the results from the cross validation for each feature set and window size combination. To produce the final results from the three runs of the cross validation, the results, i.e., the number of true positives, false positives etc. were summed to form a total for each window size and feature set combination, from the three cross validation runs (as highlighted by the number of true positives etc. in table 5.18). The totals from the three runs of the cross validation were used to calculate the overall PPV and overall sensitivity for each window size and feature set combination: the totals were used, rather than the mean across the three runs, because the latter would have resulted in non-integer values for TP, FP FN and TN, which would have been less meaningful and harder to interpret. The best results, in terms of overall PPV and overall sensitivity, from the combined results from the three iterations of the three-fold cross validation are shown in table 5.18. The window size that gave the best total results was 6 backwards 4 forwards and the following feature set: root mean square, peak to average ratio and peak to root mean square ratio.

Appliance	TP	FP	FN	TN	Sensitivity	Specificity	PPV	NPV	Overall PPV	Overall Sensitivity
Shower	24	0	0	267405	100%	100%	100%	100%	**58.26%**	**67.68%**
Microwave	21	9	9	267390	70%	99.997%	70%	99.997%		
Dishwasher	3	15	9	267402	25%	99.994%	16.67%	99.997%		
Oven	11	19	1	267398	91.67%	99.993%	36.67%	99.9996%		
Washing Machine	8	5	13	267403	38.10%	99.998%	61.54%	99.995%		

Table 5.18: Best results from the 3 iteration 3-fold cross validation (showing the totals for each of the appliances in terms of TP, etc., from the three runs of the three-fold cross validation)

Comparing the result from table 5.18 with the results in table 5.16, the results from the cross validation gave slightly worse results in terms of the overall PPV of 58.26%, compared to 62.50% in Trial 6. The overall sensitivity was lower, with the cross validation giving an overall sensitivity of 67.68%, compared with 76.92% in Trial 6.

The reason for this slight fall in overall PPV is due to the large number of false positives given by the dishwasher and the ovens, which are due to repeats in the signatures of these appliances. As highlighted in section 5.5.3.1, steps were implemented to try and alleviate the affect that these repeats have on the overall results, although the repeats of appliances have been highlighted as an issue with the recognition of appliances. This is discussed in more detail in section 5.6.4.

As cross validation has been undertaken, the results in terms of overall PPV and overall sensitivity from the trial house, as shown in table 5.18, have decreased. The cross validation was run three times, in order to produce a more reliable estimate of performance. This takes into account the variance in the results produced from these three sets of cross validation, thus giving greater confidence in the reliability of the results produced from the classifier. Due to the low usage of some of the appliances, only three fold cross validation could be carried out. As highlighted by Witten et al., (2011), this is far from optimal, and is a clear limitation, but this was all that was possible from the collected data. Due to the low number of test instances, it is not possible to calculate meaningful confidence intervals for this analysis (Witten et al., 2011).

5.6. Summary

Analysis of the data in sections 5.4 and 5.5 provide a structured approach to reviewing the full house electrical consumption data to identifying when different appliances were being used. In summary, the best results for each of the different sets of data analysis for this initial test house are shown in table 5.19. The aim of these different data models was to reach a balance between a high overall PPV and a high overall sensitivity. The best result that was achieved from this trial data analysis was by trial 6 with an overall PPV of 62.50% and an overall sensitivity of 76.92%.

Trial	Best overall PPV	Best overall sensitivity
1	1.80%	72.73%
2	3.79%	72.73%
2 with filter	20.93%	81.82%
3	5.97%	72.73%
3 with filter	20.41%	90.91%
4	3.66%	81.82%
4 with filter	18.75%	81.82%
5	28.57%	76.92%
6	62.50%	76.92%

Table 5.19: The best results from each trial

5.6.1. Window design issues

The method designed for this data analysis used a sliding window design to calculate a feature set. For trials 1-4 the window the design only “looked” forward but for the trials 5 and 6 the window design was changed to “look” forward and backwards. With this design there was an issue with other appliances being turned on or coming on within the same window as another appliance. The electric kettle in section 5.5.2.1 illustrates this problem. When the kettle was switched on there were other appliances on at the same time and this distorted the kettle’s signature and the subsequent feature set, as highlighted by figures 5.21-5.23.

It is therefore difficult for this window design method to recognise an appliance if there is another appliance on or turned on within the same window. This is a limitation of this method and the effect of it can only be limited by choosing a small window size for the analysis. The smaller the window size the lower the probability of another appliance being turned on but the choice of a smaller window size has limitations because, for this method, the window size that gave the best results was 6 forward and 4 backwards, i.e., 10 sets of 6 seconds, which equates to 60 seconds.

5.6.2. Low power appliances

The television for this house had a very low power of less than 100 Watts. This power was too low for the model to be able to recognise when the appliance was turned on, as it was lost in the background noise of the house.

The rest of the appliances that were recorded as part of this trial data analysis all had powers of greater than 1000 Watts. All of these appliances could be recognised, giving a clear signature above the background electrical activity of the house. For this house, no analysis was undertaken with appliances powered between 100 Watts and 1000 Watts. It is therefore not possible to propose the power level (between 100 Watts and 1000 Watts) where an appliance would be able to have a recognisable signature above background.

The recognition of low power appliances has been highlighted as an issue by the Franco et al., (2008), who found that it was difficult to monitor low power appliances, especially lights with a power of less than 40 Watts.

5.6.3. Appliance variability

Most of the appliances that were recorded and recognised as part of this analysis had some variability in their usage signatures due to different settings or programs chosen by the user. One example of variability in an appliance signature that is due to the actions of the user is the electric hob. The occupants were asked to note when they switched on and switched off the appliance but not the specific details about which settings and number of hobs were being used. As discussed in section 5.5.2.3, it became necessary to exclude the electric hob for this analysis, as the data that were recorded were unable to be used to train the recognition model accurately. As discussed in section 5.5.2.3, it was assumed that the reasons for these differences in signatures were due to different numbers of hobs and/or power settings of the hobs being used.

The variability of appliances signatures due to the actions of the user could cause problems when developing a model to recognise appliance usage. With many appliances there are a finite number of options; however, with some appliances there are almost an infinite number of options and it would not be possible to train the model to recognise each of these signatures.

In this household, with the exception of the hob and the washing machine, all the other appliances had repeatable signatures when used. This suggested that the user used the same settings each time the appliances were used.

The work of Franco et al., (2008), also found that the activities of the participants were regular but they were also different from each other. Therefore it is impossible to compare electricity usage for predicting appliance usage, because if one person used an appliance but another person used it with different settings or did not use that appliance does not necessarily mean that something is wrong. It just means that each person is different and everyone has different habits. However, the work of Franco et al. (2008) was undertaken with a population with the average age of 83 years, the work carried out in this trial could suggest that regularity in appliance usages and appliance settings is not just limited to elderly people.

The assumption that for this household all the same settings for all the appliances were used seems valid with the evidence. There is still the possibility that a different setting for an appliance could be used, for example, the oven being used to grill food. Without the model first being trained and tested to recognise the grill it is impossible to say if the model only trained for one setting of the oven would be able to recognise the oven if the oven was used as a grill. This could be the same as for most appliances within this house, for example, different power settings for the microwave or a different wash cycle setting for the dishwasher. Although for the washing machine there was evidence of different cycle patterns and lengths (as shown in figures 5.24 and 5.25), this did not affect the initial feature set when the appliance was turned on.

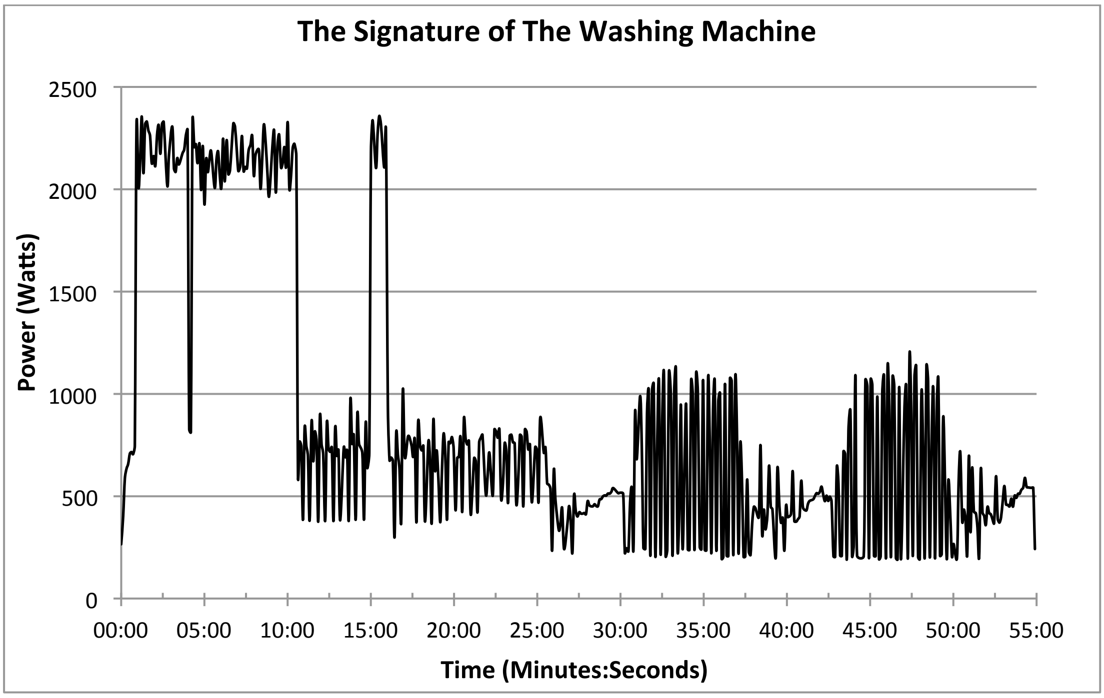

Figure 5.24: Example signature of the washing machine

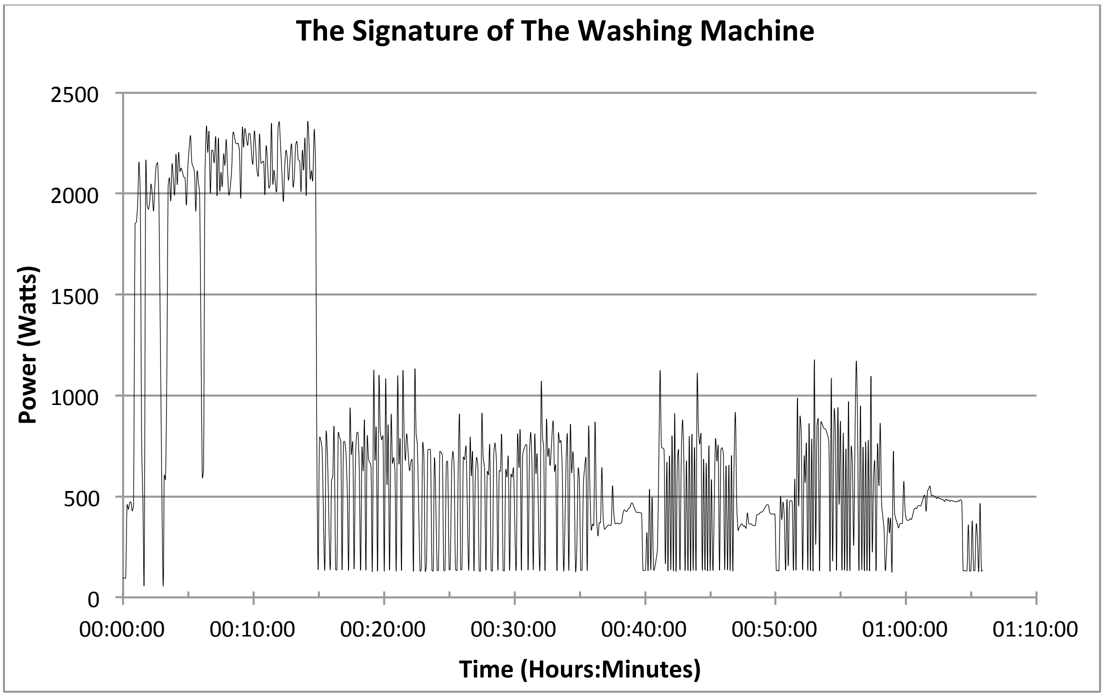

Figure 5.25: Example signature of the washing machine

It could be argued that for some appliances, for example, the washing machine, although a different wash cycle might have been used, the same process, in most cases, is still undertaken by the washing machine i.e. it will still heat up the water even if the temperature setting of the washing machine has been increased from 30 to 60. The power drawn by the washing machine for its heat up cycle cannot change only the length of time it heats up. As this model only looks at recognising the initial turn on of the washing machine, and not the length of time it has been turned on, this could explain why the model was able to recognise instances of "different" washing machine cycles.

There is evidence from this work and analysis into the different appliance signatures that, for this household, the occupants were habitual in the settings they used for appliances as well as with the appliances they used. This made it possible to develop a model for recognising different appliances, even when some of the appliances have the possibility of being variable in their signature. This work has also shown that this is not the case for every appliance, e.g., the hob, in which the variability within the appliance, due to the choice of the user, was too large to be recorded, unless the user only used the exact same settings and the same number of hobs for every use. However, this is very unlikely, although there is also the possibility that the user will change their habits, which would affect the results.

5.6.4. Appliance repeats

As shown by the figures in section 5.3, different electrical appliances have different patterns and some electrical appliances have multiple patterns for different “cycles” (for example washing machine, dishwasher, etc.)

It is these patterns that have caused a large problem within this trial data analysis with the repeats of appliances. However, all of the appliances that were recorded could have a variable length of operation time, either dependent on the user (for example, the time taken to shower) or the appliance itself (for example, how much water was in the kettle). The power used while the appliance is on generally stays “constant” in these cases and in the case of the microwave.

For other appliances, e.g., the washing machine, the dishwasher and the oven, the power usages will vary. This is shown in the example appliances signatures in section 5.3. The key with these “repeats” of signature is that they followed a similar pattern, in terms of the feature set to the feature set of when the appliance is first turned on. This means that the model recognised the repeats of certain appliances and classed them as the appliance being turned on again. This is the reason for the large number of false positives and therefore low overall PPV, highlighted by the results from trials 5 and 6 (not all the trials had all the appliances so it only affects the latter trials).

To improve the results in terms of the overall PPV, the filter was added to the model, as described in section 5.5.3.1, with the aim of “ignoring” the results of an appliance being on if there was a previous on signal from that appliance within a certain time. As shown by the results in trial 6, this filter improved the results from the model. For this filter a time limit was chosen, for each appliance, based on analysing the previous appliance patterns for the appliances in this house. This time limit may not be appropriate time limit for the same appliance from other manufacturers.

5.6.5. Conclusion

The aim of the next stage of this analysis is to repeat the method designed in this section on three further sets of electricity consumption data from three different houses. The analysis from this trial data has highlighted a number of areas that could potentially cause problems in further analysis but could also be investigated with more data.

From this analysis this method highlighted a limitation, which for future analysis could affect the results from the recognition model. This would be if there were multiple instances of other appliances being turned on or coming on when a recorded appliance is also turned on.

From this analysis there are a number of points that can be investigated further in future analysis. These are:

- **Appliance variability**: For this household, except in the case of the electric hob, possible appliance variability did not affect the ability to train and test the model. For further houses this might not be the case and could affect the ability to develop an appliance recognition model.
- **Appliance repeats**: For this household, a filter was designed to limit the affect of the “repeats” of appliances. For further households this filter might not be effective and/or might need improving. There is also a possibility that appliance repeats might not be a problem in the additional houses because different appliance manufactures might have different appliance signatures.
- **Low power appliances**: There was one low power appliance in this house, the television, which was unable to be recognised due to its power. For further households the television, if low power, or other low power appliances could cause an issue.

5.7. Conclusion

This chapter has provided an overview of the steps conducted to produce a model to recognise when an appliance, from a given list, has been used. The chapter has also provided a discussion into the issues that have been discovered from this process. In the next chapter of this thesis (Chapter 6), the method used to develop a model in this chapter will be applied to three further households, with the results and any further issues that have arisen from this analysis discussed in more detail.

Chapter 6: Electricity Data Analysis and Discussion

6.1. Introduction

Following from the development of an approach in Chapter 5, this chapter discusses the analysis and presents the results for three further sets of whole house electricity consumption data, collected from three further households. Due to the differences between the appliances, as discussed in more detail in section 6.5.7, it was not possible to transfer a single model across the houses, it was therefore decided to developed an individual model for each house based on the method that had been developed in Chapter 5. This chapter also presents a discussion of the results from all four the households and highlights any issues or problems, which have arisen through the undertaking of this analysis.

This chapter is divided into several sections. Sections 6.2, 6.3 and 6.4 present the analysis and the results for each of the three households separately, for the reason given above. Sections 6.5 and 6.6 give a discussion and a conclusion of the analysis and the results from the four households.

6.1.1. Methods

The data collection method used for the subsequent analysis shown in sections 6.2, 6.3 and 6.4, follow the same data collection method as described in section 5.2. The data was collected for a period of one week, with the actual collection dates provided in each of the sections. Ethics approval was granted from this data collection, with a copy of the ethics approval and the information sheet provided to the participant, shown in Appendix three.

6.2. Household number one

The electricity consumption data were collected between the 21st and 27th September 2013. For this household the appliances that were recorded using the diary were dishwasher, kettle, microwave oven, electric oven, television and washing machine. As this household had a gas hob a record of when the extractor fan was used was also recorded. However, after the data collection was completed the occupants for the household indicated that diary entries for the extractor fan were poorly undertaken. Because of this the extractor fan data and the diary record were excluded from the analysis.

6.2.1. Household number one- diary data

The usage data for the different appliances were extracted from the electricity consumption data using the same method as described in section 5.4.3. While conducting this data extraction an issue was discovered with one of the diary data entries for the microwave. When using the GUI (section 5.4.3) with the date and time given by the diary, there was no corresponding peak in the electricity consumption data at that time. Investigating this problem further, the electricity consumption data for that day in question showed very little activity and the participants confirmed that they were at work that day so the use of the appliance on that day and time was not possible. It was concluded by the user that they must have confused the dates. As participants could not remember the actual dates that the microwave was used and there was the possibility that more of the microwave points were recorded incorrectly, the microwave data were also excluded from further analysis.

As discussed in section 5.3.1.4, appliances using less than 100W are impossible to recognise, using single property electricity collection data. The television in this household was a more recent design and had a power rating of less than 100 Watts and so the television data was also excluded from the next stage of the analysis.

With the exclusion of the extractor fan, microwave and the television there were four appliances remaining. All of the four appliances were used frequently enough during the data collection period to provide the minimum required amount of data so that the model could be trained and tested. Although there were sufficient data for training and testing, the days used for training and testing needed to be chosen carefully as two of the four appliances were only used four times. With the model needing a minimum of two data points to train, the days for training and testing need to be chosen so as to provide sufficient data for all appliances. The training days for this household were 22nd, 23rd, 25th and 27th September and the test days were 21st, 24th and 26th September. For reference, the distribution of the appliances in terms of training and test is shown in table 6.1.

Appliance class membership	Training instances	Test instances	Total
Dishwasher	2	2	**4**
Kettle	9	9	**18**
Microwave	3	2	**5**
Oven	2	2	**4**
Television	6	10	**16**
Washing machine	3	4	**7**

Table 6.1: Distribution of training and test appliance points in Household 1

6.2.2. Household number one- trial 1 analysis

The trial followed the method used in section 5.5.4 and featured the same window design as shown in that section. This trial included all the appliances that had sufficient data for training and testing (as discussed in section 6.2.1). The training and test data for appliances were gathered using the same method as all of the previous appliances (as described in section 5.4.3) and the training and testing days were those given in section 6.2.1. As in the previous trial (section 5.5.4), this window design was run for multiple window sizes (back 1-6 and forwards 2-6) as well as different combinations of the window design's feature set (table A5.1, appendix five).

Unlike the final method as shown in section 5.5.4, this method did not include a filter to remove the repeats of certain appliances, for example, the oven, as discussed in section 5.5.3.1. This household had an oven and dishwasher and washing machine, it was decided initially to run the method without a filter to investigate whether the repeats of these three appliances affected the results in a similar way to those for the trial data analysis in Chapter 5.

6.2.3. Household number one- trial 1 results and discussion

The complete set of results from trial one are not shown in the appendix of this thesis as each run of this method for the multiple window sizes (back 1-6 and forwards 2-6) as well as different feature set combinations, produced 2970 sets of results. This would take up too many pages, and too much detail, to incorporate into the thesis and so only the final result will be shown (although the results for these trials are available on request from the author).

From analysing the results given by all the window sizes and feature set combinations, the window size and combination that gave the most acceptable results in terms of overall positive predictive value (PPV) was backwards 5 and forwards 2 (using the window design method illustrated in section 5.5.1), utilising the following features: standard deviation, root mean square and peak to root mean square ratio. The results for this best result combination, shown in table 6.2, are given in terms of true positives (TP), false positives (FP), false negatives (FN), true negatives (TN), sensitivity, specificity, positive predicted value (PPV) and negative predicted value (NPV) for each appliance. The overall PPV and overall sensitivity from all the appliances is also shown, calculated using the equation in section 5.4.6.6.

Appliance	TP	FP	FN	TN	Sensitivity	Specificity	PPV	NPV	Overall PPV	Overall Sensitivity
Dishwasher	1	4	1	38895	50%	99.990%	20%	99.997%	**16.28%**	**82.35%**
Oven	2	36	0	38863	100%	99.907%	5.26%	100%		
Kettle	8	21	1	38871	88.89%	99.946%	27.59%	99.997%		
Washing machine	3	11	1	38886	75%	99.972%	21.43%	99.997%		

Table 6.2: Results from window size 5 backwards 2 forwards, feature set (using the standard deviation, root mean square and peak to root mean square ratio)

As the results in table 6.2 show, the best overall PPV for this trial across all window sizes and feature set combinations was 16.28% with an overall sensitivity of 82.35%. In terms of overall PPV, this result does not produce a satisfactory result. The reason for this low value of PPV is the large number of false positives that are produce from all the appliances, i.e., the model predicted that the appliance was being used when it was not. Analysing all the results from this trial, the oven, in particular, produced a large number of false positives for each window size and feature set combination iteration. It was concluded from this analysis that the oven model was also picking up the “repeats” of the oven. This issue was discussed previously in more detail in section 5.5.3.1.

To reduce the number of repeats of the oven, the same filter as discussed in section 5.5.3.1.1 was added to the code for the next trial. This filter differed slightly in the time specified for the removal of the repeats of the oven than the filter discussed in section 5.6.3.1.1. Although this oven had an almost identical pattern to the example oven signature, shown in figure 5.5, its warm up time for when the oven was first turned on was found to be much longer than the oven in the household in Chapter 5. A reason for this longer time could be explained by the user setting the oven to a higher temperature than that in the example figure 5.5. Another reason for this change is a difference in oven functions between different manufacturers of the same appliance. To compensate for this change in time, the time in which the model would ignore an oven “repeat” was extended from the 15 minutes (used in section 5.5.3.1.1) to 20 minutes.

The dishwasher and washing machine also produced repeats of a similar size to when they were first turned on. The washing machine followed a similar pattern to the example given in figure 5.7, but the size of the peaks later on in the washing machine cycle were found to be much higher and similar in size to when the washing machine was first turned on. To compensate for the different signature of the washing machine, the time in which the repeats of washing machine would be ignored by the model was set to 60 minutes. The dishwasher for this household produced a similar pattern to figure 5.8 in terms of structure of the peaks but produced a different pattern in terms of number of peaks and also time. The dishwasher was found to be on for a much longer time than the example given in figure 5.8. To compensate for this difference in time, the time in which the repeats of the dishwasher would be ignored by the model was set to 40 minutes.

6.2.4. Household one- trial 2

For this trial, the method used in trial 1 was then incorporated with the filters for the repeats of the appliances, as discussed in section 6.2.3. The aim of this trial was to investigate whether the addition of a filter on the repeats of certain appliances would improve the overall PPV results. The results from this trial are discussed in the next section (6.2.5).

6.2.5. Household one- trial 2 results and discussion

As with the results from Chapter 5, this method was run for multiple window sizes (back 1-6 and forwards 2-6) as well as different combinations of the window design’s feature set. The list of the different feature set combinations that were run are shown in table A5.1 of appendix five. However, running the data for window size back 1-6 and forwards 2-6 and for the different combinations, produced 2970 sets of results, so only the best results from this trial are shown. The full results from this trial are available on request from the author.

Analysing the results from this trial, the window size and combinations that gave the best results in terms of overall PPV was backward 2 and forwards 2, using the feature set of average, standard deviation, root mean square and peak to root mean square ratio. The results for this combination are shown in more detail in table 6.3.

Appliance	TP	FP	FN	TN	Sensitivity	Specificity	PPV	NPV	Overall PPV	Overall Sensitivity
Dishwasher	1	2	1	38309	50%	99.995%	33.33%	99.997%	**54.17%**	**76.47%**
Oven	2	2	0	38309	100%	99.995%	50%	100%		
Kettle	7	5	2	38299	77.78%	99.987%	58.33%	99.995%		
Washing machine	3	2	1	38307	75%	99.995%	60%	99.997%		

Table 6.3: Results from window size 2 backwards 2 forwards, feature set (average, standard deviation, root mean square and peak to root mean square ratio)

As shown by the results in table 6.3, the best overall PPV was 54.17% with an overall sensitivity of 76.47%. Comparing the results from this trial with the results from the first trial, the results the overall PPV has increased from 16.28% to 54.17% and the overall sensitivity has decreased from 83.25% to 76.47% but the results could still be considered as acceptable. The addition of a filter to remove the repeats of the oven, dishwasher and washing machine markedly improved the overall PPV. It was evident that the filter had been successful in reducing false positives, for example, the number of false positives for the oven reduced from 36 to 2.

For reference, the attribute values for each class are provided in appendix six, table A6.2 for the best window size and feature set combination (as shown in table 6.3). As previously discussed, because the values of the attributes changes during each iteration, it would be too much information to show each attribute value for each class for each iteration.

6.2.5.1. Household number one- cross validation

Cross validation was conducted using the k-fold cross validation method, i.e., the same that was used for that described in section 5.5.6. The best total results, in terms of overall PPV and overall sensitivity, from the results from each of the three runs of the three-fold cross validation is shown in table 6.4: the total results, rather than the mean, were used for the reason explained in 5.5.6. The window size that gave the best cumulative result was: 1 backwards, 2 forwards and the feature set was: average, peak, root mean square and peak to average ratio.

Appliance	TP	FP	FN	TN	Sensitivity	Specificity	PPV	NPV	Overall PPV	Overall Sensitivity
Dishwasher	6	24	6	253065	50%	100%	20%	100%	**50.39%**	**65.66%**
Oven	7	3	5	253086	58%	99.999%	70%	99.998%		
Kettle	40	23	14	253024	74%	99.991%	63.49%	99.994%		
Washing machine	12	14	9	253066	57.14%	99.994%	46.15%	99.9964%		

Table 6.4: Best results from the 3 iteration 3-fold cross validation (showing the totals for each of the appliances in terms of TP etc. from the three runs of the three-fold cross validation)

Comparing the results from table 6.4 with the results in table 6.3, the results from the cross validation gave slightly poorer results in terms of overall PPV of 50.39%, compared to 54.17% from Trial 2, as described section 6.2.5.1. The overall sensitivity also fell with the cross validation giving an overall sensitivity of 65.56% compared with 76.47%. The reason for this slight decline in the overall PPV is due to the large number of false positives given by the dishwasher, with an individual appliance PPV of 20% and the washing machine, with an individual PPV of 46.25%. The large number of false positives for these appliances are due to repeats in their signatures. As highlighted in section 5.5.3.1 and 6.2.4, steps were taken to try and alleviate the affect these repeats on the overall results, although the repeats of appliances have been highlighted as an issue with the recognition of appliances and is discussed in more detail in section 6.5.5.

6.3. Household number two

For this household the electricity consumption data was recorded from the 22nd and 28th November 2013. For this household the appliances that were recorded using the diaries were dishwasher, kettle, microwave, television, toaster and washing machine. However, this household used gas for cooking and the user did not record any instances of the use of the extractor fan.

6.3.1. Household number two- diary data

The appliance usage data was extracted from the electricity consumption data using the same method as described in section 5.4.3.

During the data extraction it was identified that the washing machine had only be used on one day over the data collection period. Although there were enough instances of usage to be able to train the recognition model the design of this method did not allow for the splitting of days, because the model was designed to train and be tested on full days of data rather than individual instances of usage. With this design, the data could be used to train the recognition model to recognise when the washing machine was turned on but there would be no instances of washing machine usage to test. For this reason the data from the washing machine were excluded from this trial.

The television for this household was also excluded from this trial because it had a power of less than 100 Watts; the reasons for this were discussed in more detail in section 5.3.1.4.

With the exclusion of the washing machine and the television, the remaining four appliances had sufficient usage for the data to be split into training and test datasets. As discussed in section 6.2.1, the selection of the days used for training and testing had to be undertaken carefully due to the low usage of some of the appliances in this household. The training days for this household were the 23nd, 24rd, 25th and 28th November and the test days were the 22st, 26th and 27th November. For reference, the distribution of the appliances in terms of training and test is shown in table 6.5.

Appliance class membership	Training instances	Test instances	Total
Dishwasher	2	1	3
Kettle	24	23	47
Microwave	4	2	6
Toaster	2	1	3
Television	4	5	9
Washing machine	3	0	3

Table 6.5: Distribution of training and test appliance points in Household 2

6.3.2. Household number two- trial 1

The method for this trial followed the method used in section 5.5.4 (also previously in this chapter, section 6.2.4) and featured the same window design as shown in those sections. This trial included all of the appliances that had sufficient data for training and testing (as discussed in section 6.3.1). The training and test data for appliances was gathered using the same method as all of the previous appliances (as described in section 5.4.3) and the training and testing days are those given in section (6.3.1). As in the previous trial (section 5.5.4), this window design was run for multiple window sizes (back 1-6 and forwards 2-6) as well as different combinations of the window design's feature set.

The trial for this household also included a filter to remove the repeats of the dishwasher, which follows the same design as the one described in section 6.2.3. The time at which the repeats of the dishwasher would be disregarded by the model was set up to 40 minutes.

6.3.3. Household number two- trial 1 results and discussion

The results from this trial are given in the same format as those for all the previous trials, as discussed in Chapter 5 and Chapter 6. The list of the different feature set combinations that were run is shown in table A5.1 of appendix five, although running the data for window size back 1-6 and forwards 2-6 and for the different combinations, produced 2970 sets of results, so only the best results from this trial will be shown. The full results from this trial are available on request from the author.

From analysing all of the results given by the different window sizes and feature sets, the window size and feature set combination that gave the best results in terms of overall PPV was window size 1 backwards and 3 forwards, using a feature set of standard deviation, peak to average ratio and root mean square to average ratio. The results for this are shown in more detail in table 6.6.

Appliance	TP	FP	FN	TN	Sensitivity	Specificity	PPV	NPV	Overall PPV	Overall Sensitivity
Toaster	0	0	1	38619	0%	100%	NaN[6]	99.997%	**90.48%**	**70.37%**
Kettle	19	2	4	38595	83%	99.995%	90.48%	100%		
Dishwasher	0	0	1	38619	0%	100%	NaN	99.997%		
Microwave	0	0	2	38618	0%	100%	NaN	99.995%		

Table 6.6: Results from window size1 backwards 3 forwards, feature set (standard deviation, peak to average ratio and root mean square to average ratio)

From the results in table 6.6, the overall PPV was 90.48% with the overall sensitivity of 70.37%. However, looking at the results in terms of true positives, false positives etc. the model only recognised instances of the kettle for this window size and

[6] NaN is used, as it is not possible to define zero divided by zero.

feature set combination. This is highlighted by the other rows giving no results (zeros), in terms of true positives and false positives. Although, unlike other instances where the best overall PPV has been given by window size and feature set combination where the model has not recognised an instance of usage for all appliances, the overall sensitivity for this window size combination is much higher. In fact, the overall sensitivity of this window size combination is similar to those given in 6.2.5 and 5.6 as the best results from each respective household.

The reason for these high overall PPV and sensitivity may be the low number of times of usage of the other appliances. For this household, two of the other four appliances that were recorded were only used three times which, after data required for training, left only one instance for testing. Although it could be argued that the results, just in terms of value for overall PPV and sensitivity would be satisfactory, the results from this trial were analysed again to see whether there was window size and feature set combination that gave satisfactory results in terms of PPV, sensitivity and recognising at least one instance of use (true positive) for each appliance. The window size and feature set combination that gave the best results where all three of the criteria were met was a window size of 2 backwards and 6 forwards, using a feature set of average, peak, standard deviation, root mean square, peak to average ratio and peak to root mean square ratio. The results for this are shown in table 6.7.

Appliance	TP	FP	FN	TN	Sensitivity	Specificity	PPV	NPV	Overall PPV	Overall Sensitivity
Toaster	1	4	0	38920	100%	99.990%	20%	100%		
Kettle	20	3	3	38899	86.957%	99.992%	86.957%	99.992%	**72.73%**	**88.89%**
Dishwasher	1	1	0	38923	100%	99.997%	50%	100%		
Microwave	2	1	0	38922	100%	99.997%	66.667%	100%		

Table 6.7: Results from window size 2 backwards 6 forwards, feature set (average, peak, standard deviation, root mean square, peak to average ratio and peak to root mean square ratio)

From table 6.7, it can be seen that the result that meets the criteria of a high overall PPV, sensitivity and recognising at least one instance for each appliance was an overall PPV of 72.73% and an overall sensitivity of 88.89%. However, for this household there was a large number of window sizes and feature set combinations that give a higher value for overall PPV with the same value for overall sensitivity, because this window size and feature set combination did not provide at least one true positive for all appliances.

If there was the possibility of collecting further data for this household, it is possible that the best result in terms of overall PPV and overall sensitivity could be improved. Due to the single usage of the toaster and the dishwasher in this household, over the data period collection, these appliances could not have any false negatives. This is different to the other households in this study where the appliances could have false negative predictions, as long as there were at least one instances of appliance usage recognition.

For this household, there were window size and feature set combinations that gave a high recognition rate for the kettle, with 20 true positives, three false positives and three false negative. Although for this household, especially with its low amount of usage data, this method is a compromise in finding the best feature set and window size combination that represents all the appliances and, in some cases, better recognition results for individual appliances had to be sacrificed for a satisfactory recognition results for all appliances.

For reference, the attribute values for each class are provided in appendix six, table A6.3, for the best window size and feature set combination (as shown in table 6.7). As previously discussed, because the values of the attributes changes during each iteration, it would be too much information to show each attribute value for each class for each iteration.

6.3.3.1. Household number two- cross validation

Cross validation was undertaken using the k-fold cross validation method, which is the same as that used as previously used in section 5.5.6. The best total results in terms of overall PPV and overall sensitivity from the results from each of the 3 runs of the 3-fold cross validation is shown in table 6.8 (the total results were used for the reason explained in 5.5.6). The window size that gave the best cumulative result was 3 backwards 5 forwards and the following feature set: average, peak, standard deviation and root mean square to average ratio.

Appliance	TP	FP	FN	TN	Sensitivity	Specificity	PPV	NPV	Overall PPV	Overall Sensitivity
Toaster	3	5	6	271702	33%	100%	38%	100%	**61.20%**	**63.28%**
Kettle	97	34	44	271541	69%	99.987%	74%	99.984%		
Dishwasher	6	26	3	271681	67%	99.990%	18.75%	99.999%		
Microwave	6	6	12	271692	33.33%	99.998%	50.00%	99.996%		

Table 6.8: Best results from the 3 iteration 3-fold cross validation (showing the totals for each of the appliances in terms of TP etc. from the three runs of the three fold cross validation)

Comparing the results from table 6.8 with the results in table 6.7, the results from the cross-validation gave a poorer result in terms of overall PPV of 61.20% compared to 72.73%. The overall sensitivity also declined with the cross validation giving an overall sensitivity of 63.28%, compared with 88.89% in the Household 2, Trial 1 model. The reason for this decline in overall PPV is due to the large number of false positives given by the dishwasher and the kettle, with an individual appliance PPV for the dishwasher of 18.75%.

Other window size and features set did provide better results in terms of overall PPV and sensitivity although this did not meet the requirement of recognising at least one instance of correct classification for each of the appliances. Due to the limited number of data points, the maximum number of folds that could be used was 3. As highlighted by Witten et al., (2011), for optimal results, 10-fold cross validation should be used, this is addressed in the recommendations made for future work in section 7.5.

6.4. Household number three

The electricity consumption data were collected from the dates of the 13st to the 20th October 2013. For this household the appliances that were recorded using the diary were shower, kettle, toaster, television and washing machine. As this household also had a gas hob the household were also asked to record when the extractor fan was used.

6.4.1. Household number three- diary data

The appliance usage data was extracted from the electricity consumption data using the same method as described in section 5.4.3. It was during the data extraction that the toaster data was split into two appliances for recognition. The reason for this was that the toaster for this household could be used for four or two pieces of bread. The user had recorded in the diary whether the appliance had been turned on in the four or the two-piece mode. In looking at the differences in the electrical data when the toaster was used in the different modes, as might be expected the power consumption of the toaster was almost double in the four-piece mode. As the appliance was acting in two power ranges it was decided to treat the results as being from two appliances.

As discussed in section 5.3.1.4 the television for this household was excluded from the next stage, as its power was less than 100 Watts. Although the aim had been to include the extractor fan as a proxy to indicate cooking, the extractor fan for this household was also found to be less than 100 Watts and was impossible to distinguish within the noise of the data and so was also excluded from the analysis.

With the exclusion of the extractor fan and the television, the remaining four appliances had sufficient usage for the data to be split into training and test datasets. As discussed in section 6.2.1, the selection of the days used for training and testing had to be made carefully due to the low usage of some of the appliances in this household. The training days for this household were the 13th, 14th, 17th, 18th and 20th October and the test days were the 15th, 16th and 19th October. For this household the data was collected over 7 days but two of the days (13th and 20th) only contained half days of data. For reference, the distribution of the appliances in terms of training and test is shown in table 6.9.

Appliance class membership	Training instances	Test instances	Total
Extractor fan	5	4	**9**
Kettle	38	28	**66**
Shower	6	5	**11**
Toaster	12	7	**19**
Television	8	10	**18**
Washing machine	5	3	**8**

Table 6.9: Distribution of training and test appliance points in household 3

6.4.2. Household number three- trial 1

The method for this trial followed the method used in section 5.5.4 (also outlined previously in this chapter, section 6.2.4) and featured the same window design as shown in that section. This trial included all the appliances, which had sufficient data for training and testing (as discussed in section 6.4.1). The training and test data for appliances was gathered using the same method as all of the previous appliances (as described in section 5.4.3) and the training and testing days are those given in section (6.4.1). As in the previous trial (section 5.5.4), this window design was run for multiple window sizes (back 1-6 and forwards 2-6) as well as different combinations of the window design's feature set.

The trial for this household also included a filter to remove the repeats of the washing machine, which followed the same design as the one described in section 6.2.3. The time at which the repeats of the washing machine would be discarded was set up to 40 minutes.

6.4.3. Household number three- trial 1 results and discussion

The results from this trial are given in the same format as those for all the previous trials, as discussed in Chapter 5 and earlier in Chapter 6. The list of the different feature set combinations that were run is shown in table A5.1 of appendix five. For reasons described previously, only the best results from this trial are shown. The full results from this trial are available on request from the author.

From analysing all the results given by the different window sizes and feature sets, the window size and feature set combinations, there were four window combinations that gave the best results in terms of overall PPV. All four of these results, and the details of respective window size and feature set combinations are shown in more detail in tables 6.10 – 6.13.

Appliance	TP	FP	FN	TN	Sensitivity	Specificity	PPV	NPV	Overall PPV	Overall Sensitivity
Shower	5	0	0	38556	100%	100%	100%	100%	**100%**	**83.72%**
Washing machine	2	0	1	38558	66.67%	100%	100%	99.997%		
Kettle	25	0	3	38533	89.29%	100%	100%	99.992%		
Toaster (2)	3	0	2	38556	60%	100%	100%	99.995%		
Toaster (4)	1	0	1	38559	50%	100%	100%	99.997%		

Table 6.10: Results from window size 3 backwards 4 forwards, feature set (root mean square, peak to average ratio and peak to root mean square ratio)

Appliance	TP	FP	FN	TN	Sensitivity	Specificity	PPV	NPV	Overall PPV	Overall Sensitivity
Shower	4	0	1	38556	80%	100%	100%	99.997%	**100%**	**81.40%**
Washing machine	2	0	1	38558	66.67%	100%	100%	99.997%		
Kettle	25	0	3	38533	89.29%	100%	100%	99.992%		
Toaster (2)	3	0	2	38556	60%	100%	100%	99.995%		
Toaster (4)	1	0	1	38559	50%	100%	100%	99.997%		

Table 6.11: Results from window size 3 backwards 4 forwards, feature set (root mean square, peak to root mean square ratio and root mean square to average ratio)

Appliance	TP	FP	FN	TN	Sensitivity	Specificity	PPV	NPV	Overall PPV	Overall Sensitivity
Shower	5	0	0	38556	100%	100%	100%	100%	**100%**	**79.07%**
Washing machine	2	0	1	38558	66.67%	100%	100%	99.997%		
Kettle	23	0	5	38533	82.14%	100%	100%	99.987%		
Toaster (2)	3	0	2	38556	60%	100%	100%	99.995%		
Toaster (4)	1	0	1	38559	50%	100%	100%	99.997%		

Table 6.12: Results from window size 3 backwards 4 forwards, feature set (peak, peak to average ratio and peak to root mean square ratio)

Appliance	TP	FP	FN	TN	Sensitivity	Specificity	PPV	NPV	Overall PPV	Overall Sensitivity
Shower	4	0	1	38556	80%	100%	100%	99.997%	**100%**	**79.07%**
Washing machine	2	0	1	38558	66.67%	100%	100%	99.997%		
Kettle	24	0	4	38533	85.71%	100%	100%	99.990%		
Toaster (2)	3	0	2	38556	60%	100%	100%	99.995%		
Toaster (4)	1	0	1	38559	50%	100%	100%	99.997%		

Table 6.13: Results from window size 3 backwards 4 forwards, feature set (peak, peak to root mean square ratio and root mean square to average ratio)

As the results from tables 6.10 to 6.13 show, the best overall PPV of 100% with the best overall sensitivity varying from 83.72% to 79.07%. Although all the results in tables 6.10 to 6.13 give the same results in terms of overall PPV, there are differences in the numbers of true positives and false negatives for some of the appliances between the different window size combinations. An example of this is the difference between the number of true positives for the shower in tables 6.10 and 6.11, with the shower in table 6.10 having 5 true positives and no false negatives and the shower in table 6.11 having 4 true positives and 1 false negative.

Although the results shown in the four tables above could all be classed as acceptable, there also needs to be a balance between achieving the best overall

PPV and the best overall sensitivity. The results were analysed again in terms of improving the best overall sensitivity and the results are discussed below.

The best overall sensitivity that was given by this method was 90.70%, this result was given by multiple window sizes and feature set combinations. The highest value of overall PPV that is given with this sensitivity is 95.12%. There are five window size and feature set combinations that gave these results, as three of these window sizes give the same results in terms of true positives, true negatives, false positives false negatives for all the appliances. The results for these window size and feature set combinations are shown in tables 6.14 to 6.16.

Appliance	TP	FP	FN	TN	Sensitivity	Specificity	PPV	NPV	Overall PPV	Overall Sensitivity
Shower	5	0	0	37796	100%	100%	100%	100%	**95.12%**	**90.70%**
Washing machine	3	1	0	37797	100%	99.997%	75%	100%		
Kettle	25	0	3	37773	89.286%	100%	100%	99.992%		
Toaster (2)	5	1	0	37795	100%	99.997%	83.333%	100%		
Toaster (4)	1	0	1	37799	50%	100%	100%	99.997%		

Table 6.14: Results from window size 1 backwards 4 forwards, feature set (average, peak, standard deviation, root mean square and root mean square to average ratio)

Appliance	TP	FP	FN	TN	Sensitivity	Specificity	PPV	NPV	Overall PPV	Overall Sensitivity
Shower	5	0	0	38556	100%	100%	100%	100%	**95.12%**	**90.70%**
Washing machine	3	1	0	38557	100%	99.997%	75%	100%		
Kettle	26	1	2	38532	92.86%	99.997%	96.296%	99.995%		
Toaster (2)	4	0	1	38556	80%	100%	100%	99.997%		
Toaster (4)	1	0	1	38559	50%	100%	100%	99.997%		

Table 6.15: Results from window size 3 backwards 4 forwards, feature set (peak, standard deviation, root mean square and root mean square to average ratio)

Appliance	TP	FP	FN	TN	Sensitivity	Specificity	PPV	NPV	Overall PPV	Overall Sensitivity
Shower	5	0	0	38243	100%	100%	100%	100%		
Washing machine	3	1	0	38244	100%	99.997%	75%	100%		
Kettle	26	0	2	38220	92.86%	100%	100%	99.995%	**95.12%**	**90.70%**
Toaster (2)	4	1	1	38242	80%	99.997%	80%	99.997%		
Toaster (4)	1	0	1	38246	50%	100%	100%	99.997%		

Table 6.16: Results from window size 3 backwards 3 forwards, feature set (standard deviation, root mean square and root mean square to average ratio.) Results from window size 5 backwards 2 forwards, feature set (peak, standard deviation, peak to average ratio and peak to root mean square ratio.) Results from window size 5 backwards 2 forwards, feature set (average, peak, standard deviation, root mean square, peak to average ratio and peak to root mean square ratio.)

As shown by the results in tables 6.14 to 6.16, all the different window size and feature set combinations give the same results of an overall PPV of 95.12% and an overall sensitivity of 90.70%. Although the results in Tables 6.14 to 6.16 all the gave the same overall results, each of the tables has a slightly different set of results in terms of true positives etc. for some of the appliances. This is because each of the results in tables 6.14 to 6.16 represents a compromise between the overall results and the results for each of the appliances and the number of true positives, false positives, false negatives and true negatives they each have.

Although both of the overall results shown in tables 6.10 to 6.13 and tables 6.14 to 6.16 would be acceptable, there needed to be a trade-off between the overall PPV and overall sensitivity. Even though the results in tables 6.14 to 6.16 produce a lower PPV of 95.12% than in tables 6.10 to 6.13 (100%), the overall sensitivity is also much higher, i.e., 90.70% compared with that in tables 6.10 to 6.13 (83.72%). Therefore, the results in tables 6.14 to 6.16 produce a much better trade off between the values of overall PPV and sensitivity.

For reference, the attribute values for each class are provided in appendix six, table A6.4 for the best window size and feature set combination (as shown in table 6.15). As previously discussed, because the values of the attributes change through each iteration, it would be too much information to show each attribute value for each class for each iteration.

6.4.3.1. Household number three- cross validation

Cross validation was conducted with the k-fold cross validation method used which was the same as that previously used, as shown in section 5.5.6. The best total results in terms of overall PPV and overall sensitivity from the results from each of the 3 runs of the 3-fold cross validation is shown in table 6.17 (the total results were used for the reason explained in 5.5.6). The window size that gave the best cumulative result was 5 backwards 4 forwards and the following feature set: standard deviation, peak to root mean square ratio and root mean square to average ratio.

Appliance	TP	FP	FN	TN	Sensitivity	Specificity	PPV	NPV	Overall PPV	Overall Sensitivity
Shower	32	0	1	278268	97%	100%	100%	100%	**86.34%**	**89.10%**
Washing machine	19	5	5	278272	79%	99.998%	79%	99.998%		
Kettle	185	6	13	278097	93%	99.998%	96.86%	99.995%		
Toaster (2)	33	30	12	278226	73.33%	99.989%	52.38%	99.9957%		
Toaster (4)	9	3	3	278286	75.00%	99.999%	75.00%	99.999%		

Table 6.17: Best results from the 3 iteration 3-fold cross validation (showing the totals for each of the appliances in terms of TP etc. from the three runs of the three-fold cross validation)

Comparing the results in table 6.17 with the results in tables 6.14-6.16, the results from the cross validation gave a poorer result in terms of overall PPV of 86.34% compared to 95.12%. The overall sensitivity also declined slightly with the cross validation giving an overall sensitivity of 89.10% compared with 90.70%. Another window size and feature sets did provide a better result in terms of overall PPV, with an overall PPV of 89.50%, although it gave a worse sensitivity of 79.81%. As discussed in section 6.4.3, the choice of the best window is a trade off between the values of overall PPV and overall sensitivity.

6.5. Discussion

6.5.1. Introduction

This section provides a discussion of the results of the analyses of the four households, undertaken in sections 6.2, 6.3, 6.4 and 5.5.4. This section will place this work into the context of other current work in this area.

The aims of this research, as highlighted in section 1.6, were to examine the feasibility of collection of this type of electricity data from a number of households, to analysis this data to highlight in the different settings when appliances and thus, by inference, activities that had taken place. In addition, the aim was to make recommendations on the use of a single whole household electricity consumption monitor as a method for activity recognition. This section will address these aims, as well as to address further points of discussion that have arisen from the analysis of this electricity data.

This section is split into sections with section 6.5.2 giving an overview of the best results from each of the four households. Section 6.5.3 then discusses certain limitations with the window design and approach used from this appliance recognition. Section 6.5.4 highlights some of the issue with recognising low power appliances. Section 6.5.5 discusses the problems with appliance repeats. Section 6.5.6 highlights the variability in appliances. Finally, section 6.5.7 will show the differences in electricity consumption across multiple households.

6.5.2. Overview of results

As shown by the results in sections 6.2, 6.3, 6.4 and 5.5.4, four sets of electricity consumptions data were analysed with the best results from each of theses households summarised in table 6.18.

Household	Feature set	Best window size	Best overall PPV	Best overall sensitivity
Trial household	Peak Standard deviation Peak to average ratio	6 backwards, 4 forwards	62.50%	76.92%
	Standard deviation Root mean square Peak to root mean square ratio	6 backwards, 4 forwards	62.50%	76.92%
Household one	Average Standard deviation Root mean square Peak to root mean square ratio	2 backwards, 2 forwards	54.17%	76.47%
Household two	Average Peak Standard deviation Root mean square Peak to average ratio Peak to root mean square ratio	2 backwards, 6 forwards	72.73%	88.89%
Household three	Average Peak Standard deviation Root mean square Root mean square to average ratio	1 backwards, 4 forwards	95.12%	90.70%
	Peak Standard deviation Root mean square Root mean square to average ratio	3 backwards, 4 forwards	95.12%	90.70%
	Standard deviation Root mean square Root mean square to average ratio	3 backwards, 3 forwards	95.12%	90.70%
	Peak Standard deviation Peak to average ratio Peak to root mean square ratio	5 backwards, 2 forwards	95.12%	90.70%
	Average Peak Standard deviation Root mean square Peak to average ratio Peak to root mean square ratio	5 backwards, 2 forwards	95.12%	90.70%

Table 6.18: Table of the best results from each trial

Table 6.18, gives a summary of the best results from each of the four households. As described in section 5.4.6.6, the criteria for choosing the best results from each of the models were the overall PPV and overall sensitivity values. These were chosen, as the aim of this model was to maximise the number of positive identifications of the appliances being switched on from all the predictions it made of the appliance being switched on, i.e. high overall PPV, and to also recognise, with a high level of accuracy, when an appliance had not been used (i.e. high overall sensitivity). The choice of the best results from some of these households has been a compromise between these two values, as discussed in section 6.4.3.

For the trial household and household number three, a number of multiple window sizes and feature set combinations gave the same results, in terms of overall PPV and overall sensitivity. However, some of these combinations did give different results in terms of true positives, false positives, false negatives and true negatives for the appliances (shown in the tables of results in sections 6.4.3 and 5.5.5).

Table 6.18 also highlights that there was not a common window size and feature set combination that would give the best result across all of the households. This highlights the differences across the patterns of household electricity consumption data from multiple households; these differences are discussed in more detail in section 6.5.7.

Having analysed the results from each of the four households, it is understandable that each household would produce a different window size and feature set combination that gave the best results, because four different models were created, having been trained only on the appliance data from that household. This also shows that the pattern of electricity consumption data is unique to the household (and its occupants and possibly other variables, e.g., time of the year), even though the same appliances (if present in the household) were recorded in each of the households. The results suggest that the model cannot be transferred and used on another household without first undertaking training of the data for that household. The reasons for this are discussed in more detail in section 6.5.6 of this discussion.

As shown by the results in table 6.18, the results from the four households vary widely, with the results for overall PPV varying from 54% to 95% and the results for overall sensitivity varying from 76% to 90%. However, with these variations in the results, it is not possible to say that the model developed using this approach from household one is particularly poor and the model developed from household three is particularly good, because they only reflect the different patterns of activity and electricity consumption in the different households. As they are not interchangeable, they are unique to the data of that household of which they were trained. However, from these results, it is possible to explain why the same approaches used to develop these models produced such different results. The reasons for these differences are discussed in more details in sections 6.5.7 of this discussion.

This section has given an initial overview of the results from the four households that took part in this study. It has also highlighted some areas of discussion that will be addressed in subsequent sections of this discussion.

6.5.3. Window design limitation

6.5.3.1. Window design

As discussed in the summary of the results from the trial analysis, section 5.6.1, there was a limitation with this window design. From the further analysis of data from three more households, this limitation was still present but steps were taken to limit its effect on the results. It was realised from this further data analysis, that some of the feature set data, which were used to train the model to recognise an appliance, were distorted by having an appliance turned on/off within the same window. To limit the effect of this, these points were removed from the training data, as they gave an untrue representation of the feature set of that appliance. For the instances where the feature set points had to be removed from the training data, there was still a minimum of two data points, in order to be able to meet the requirement to train the model and did not result in the exclusion of any recorded appliances from the training data.

There is also the possibility that this limitation also had an effect on the overall results of the model. This is due to another appliance being turned on/off when one of the recognised appliances was being turned on, as it would not make it possible for the model to recognise when the appliance had been turned on because the feature set would be distorted by the other appliance. From this test data from the initial trial household (Chapter 5), there were two test points for the dishwasher, although for one of these test points (as highlighted by the data in figure 6.1) there was another appliance that turned on within 6 seconds of the dishwasher being turned on. The addition of this other appliance being turned on within the same window as the dishwasher, distorted the feature set for that point. This meant that it was impossible for the model to recognise this feature set point as the dishwasher. This would, therefore, have an effect on the overall results of the model, in terms of overall PPV or overall sensitivity, as this point from the dishwasher would always be classed as a false negative (as shown by the results in table 5.16).

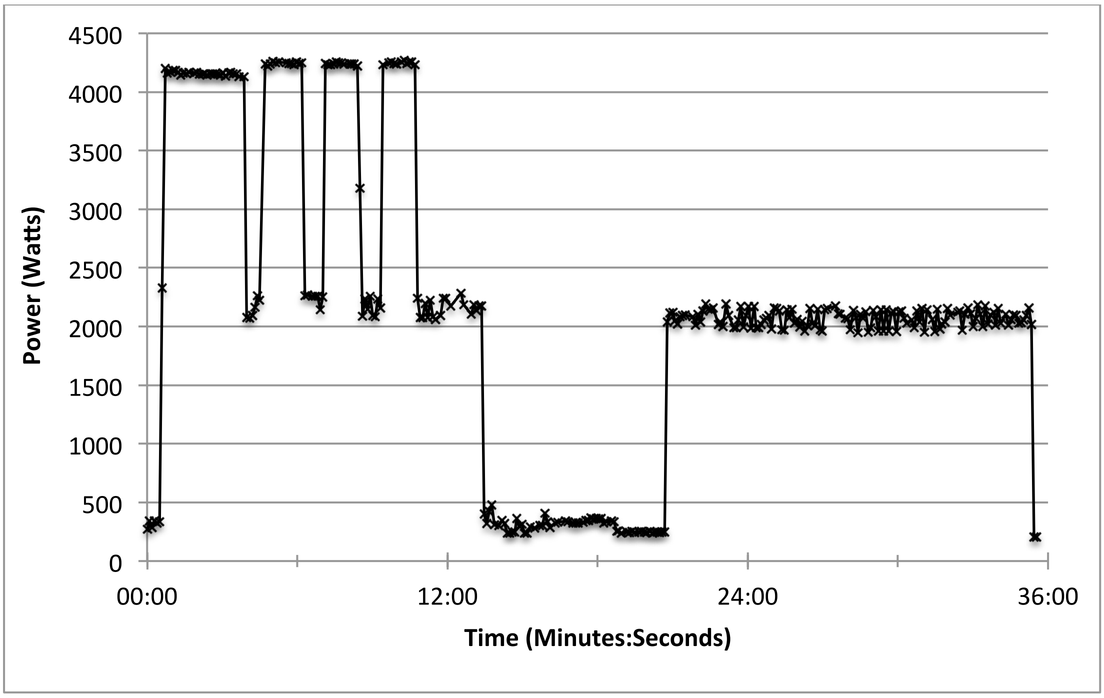

Figure 6.1: figure showing two appliances coming on within the same window

This section has highlighted how the limitation of the window design used for this approach can have an effect not only on the ability to train the model, but also on the results given from the model.

6.5.3.2. Method design

The analysis of the electricity consumption data for household two (section 6.3) highlighted an issue with the design of this approach for recognising appliances from electricity consumption data. For the test and training datasets, the data were split based on days of the week rather than just the amount of data. The rationale of this was to be able to highlight patterns of usage, and therefore the activity of an older person, throughout the days. For this approach, to be able to provide training and test data, an appliance had to be used on a minimum of two separate days. During the data collection phase, this was not highlighted to the occupants, as they were asked not to change from their usual activities for the purpose of the project.

In the case of the washing machine in household number two, there was not enough data to do training and testing as this appliance was only used a number of times on only one day of the data collection week. This is a limitation of the design of this method, although it did only affect one appliance in one household so its effect was only minimal. If the data collection period were to be extended then the used of an appliance only one day a week, for example “washing day”, would be negligible.

The discussion in this section has highlighted a limitation with the design of this approach for appliance recognition and a simple way of minimising this limitation for any future work.

6.5.4. Low power appliances

6.5.4.1. The television

As highlighted in the summary of the trial data analysis (5.6.2), there was an issue with trying to recognise low power appliances from the electricity consumption data. From the trial analysis in Chapter 5, the television was excluded from the data analysis. This was because the power drawn by the television was less than 100Watts and it was not possible to be able to be recognised over the background noise, e.g., lights, etc., and other high power appliances of the electricity consumption data. From the analysis of the further households in sections 6.2, 6.3 and 6.4, the televisions in these three households were also excluded from the data analysis as the power used by all of them was found to be less than 100Watts.

For further work in this area, it would be recommended to remove the television from the list of activities to be recorded, as the appliance cannot be reliably recognised by this method, although it would still be possible to record the use of the television, (if desired) by means of another method of monitoring. An example of this would be an individual appliance monitor, placed on the plug of the television, as discussed in Chapter 2 section 2.4.1.1. However, the addition of this type of monitor would detract slightly from this method of monitoring being non-intrusive (i.e., not visible). It would also set the conditions for the user that the television, in this case, had to be used solely by that plug and though this addition would detract from the overall aims of this approach, it would enable other appliances to be recorded, which could not reasonably be recognised from the whole household electricity consumption data.

6.5.4.2. The extractor fan

From the analysis in section 6.4, there was another appliance that was recorded as part of the data collection that had to be excluded, because its power was less than 100Watts. This was the extractor fan. The aim of recording this appliance was to enable it to be used as a proxy for food preparation in households that had gas for cooking; however, although three of the four households used gas for either their hob and/or their oven, only one household recorded their usage of the extractor fan. Because it was only recorded by one household, it is not possible to conclude if it is possible to use an extractor fan as a proxy for cooking in homes with gas or if the power of the extractor fan is too low for recognition by this type of monitoring.

Unlike the television, it would not be possible to place a plug sensor (section 2.4.1.1) on the extractor fan. This is due to the fact that these types of sensor require the appliance to have a plug and to be plugged into the mains power circuit of the household. An extractor fan is generally directly wired into the electric circuit of the household and therefore this type of monitoring is not possible.

This section has highlighted two appliances with a low power usage, which limits their ability to be recognised using this type of electricity monitor. This section has suggested ways of included some of these appliances, with the use of additional monitoring sensors, although it must be acknowledged that this is not possible for some electrical appliances within the home.

6.5.5. Appliance repeats

As described in section 5.5.3.1 and highlighted by figures 5.5, 5.7 and 5.8 in section 5.3, appliance repeats are so called because the appliance signature when it is turning itself off and on is similar to when the appliance was first turned on. These “repeats” create a problem when undertaking this type of appliance recognition, because the model recognises each of these repeats as the appliance being turned on, meaning that the model produced a large number of false positives. To address this issue a filter was added so as to remove a repeat point if there were a subsequent point within a certain specified time. The improvement in the results of the model with the addition of a filter is shown by the differences in the results of the model from trial 5 (section 5.5.2) which did not have a filter and trial 6 (section 5.5.4) which had a filter.

The repeats of some appliances also created problems of a large number of false positives for the three further households, which are analysed in this chapter. The differences in the results from the model without a filter and a model with a filter were examined in the household number one (section 6.2). As the results show in the section, the addition of a filter reduced the number of false positives from the appliances and therefore gave better results in terms of overall PPV and overall sensitivity.

The times for each of the filters were chosen based on the analysis of the signal of the appliance for each of the households. This meant that the times for the filters, for the same appliance, differed between the households. The reasons for this will be discussed in more detail in section 6.5.7, although this has highlighted some of the differences between manufacturers of the same appliance and their differences in cycle settings chosen by the user.

This section has highlighted how the signatures of some appliances and their “repeats” can cause problems and affect the overall results of the model. It has also highlighted some of the differences between manufacturers of the same appliances and their different settings. The differences in signature between the same appliance type will be discussed in more detail in section 6.5.7.

This section has also highlighted an issue with creating a generic model for appliance recognition, because the appliances need to be analysed first: the signatures of the appliances are very different and the addition of a generic filter per appliance could not effective for some appliances.

6.5.6. Appliance variability

This section will highlight some of the issues with trying to recognise the same appliance, and variability of the signature. An example of where an appliance has a variable signature arises due to the choices of the user for the electric hob in the trial household (section 5.5.2.3). Due to the variability in this appliance signature and the corresponding variability in the feature set, the electric hob was removed from the appliance recognition.

For the appliances in the three households, which were analysed in this chapter, no appliances were removed from the recognition due to their being too variable in their signatures and subsequent feature sets. However, from the analysis of these households, certain issues have arisen which highlight the challenges when trying to recognise appliances that can be variable in their signatures.

To discuss these challenges in more detail, this section is divided into two groups of appliance variability. The first of these sections will discuss appliances that have signatures that are markedly different. The second section will discuss appliances that have similarities as well as differences.

6.5.6.1. Appliance variability- different signatures

The toaster from household number three (section 6.4) is an example of an appliance that produces a markedly different signature depending on choices made by the user. For this toaster there were two settings of usage, two-piece or four-piece. Due to the differences between these two settings, the recognition of the toaster was split to form two appliances. This was possible as the user recorded when they had used the toaster either in four-piece mode or two-piece mode. In addition, the signatures of the toaster for the two and four piece modes were very different, in power usage, to each other, i.e., the four-piece mode being almost double the power of that of the two-piece mode, which was to be expected.

It was possible to be clear that the toaster had two different signatures, because the user recorded the mode in which the appliance was being used. This was an extension to what the participants were asked to do: if the user had not done this then it would have to be concluded that the toaster had two very different signatures. The presence of appliances, which can produce markedly different signatures, highlights a challenge with trying to develop a model to recognise appliance usage.

A consequence of this is that in order to be able to predict with a high degree of accuracy that the model can recognise the appliance, the model would first have to be trained on all the possible signatures of that appliance. For some appliances, the different signatures could be almost infinite, for example in the case of the electric hob. In the case of this toaster, if the appliance had only be used in two-piece mode and trained as such, the model would not recognise if the toaster were used in four-piece mode and vice versa although, that said, there were only two possible combinations for this device.

Although the presence of appliances that produce markedly different signatures can cause challenges in appliance recognition, there was also evidence that the occupants of the household were habitual in the settings they used in appliances with variable settings. It could be argued, therefore, that to train a model to recognise all instances of usage would not be necessary, because the user might not use them all.

Although, for household number three (section 6.4), the toaster produced markedly different signatures, both of the settings were used frequently by the occupants. It was, therefore, possible to train the model to recognise both types of signature and did not create a problem with the recognition. However, if the occupants were to change the settings of appliances, or how they used them, then more model training would be needed to recognise the changes.

The data collection period for this research was short with collecting only over one week, so that the burden on users for recording use was minimised. If the data collection was extended to a longer period then different appliance settings used by the occupants could be investigated further.

6.5.6.2. Appliance variability- similar signatures

From the analysis of the four households, undertaken for this research, there were also individual appliances that showed variability in their signatures, although for these appliances the difference in their signatures did not affect their ability to be recognised by the model when they were first turned on. For the washing machine in the trial household (described in section 5.6.3), there was evidence of different cycle lengths and patterns but this did not affect the initial feature set of the appliance when it was first turned on. The conclusion from these different signatures was that although the choice of the user does affect the overall signature, the washing machine would still result in the same initial signature. The washing machine could draw a finite amount of power and the examples of the two different washing machine signatures in figures 5.24 and 5.25 showed that the length of the cycles changed but the power drawn by the washing machine when it was first turned on did not.

For the washing machine in household number three (section 6.4), there was also evidence of different signatures when the washing machine was in use although, as with the washing machine signatures from the trial household in Chapter 5, the power when the washing machine was first turned on was similar. However, the overall signatures of this washing machine varied as shown by the two signatures of the washing machine from household number three in figures 6.2 and 6.3.

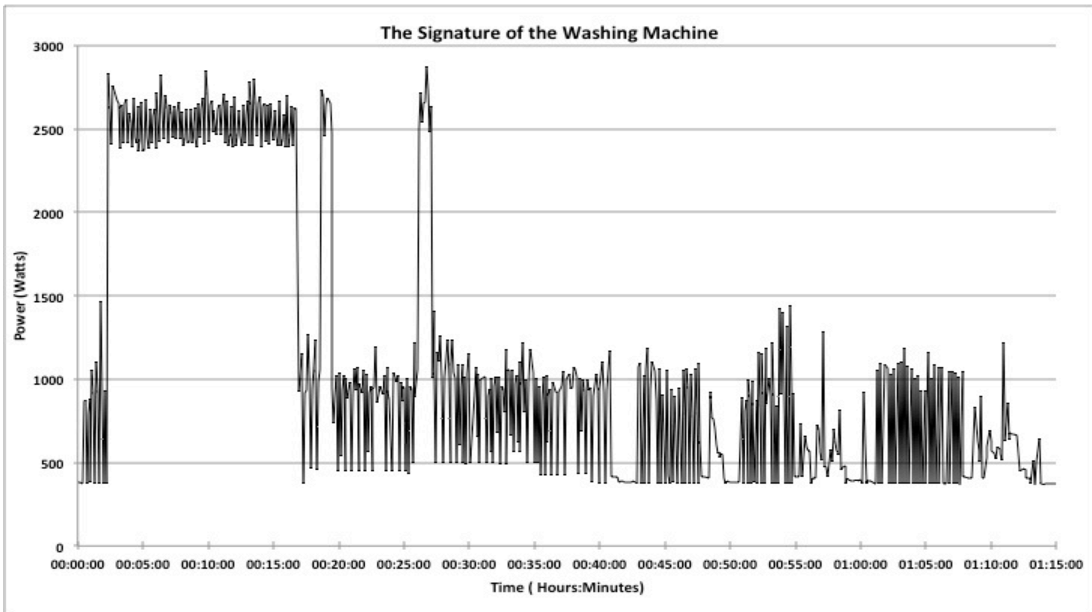

Figure 6.2: An example of the washing machine signature in the trial household

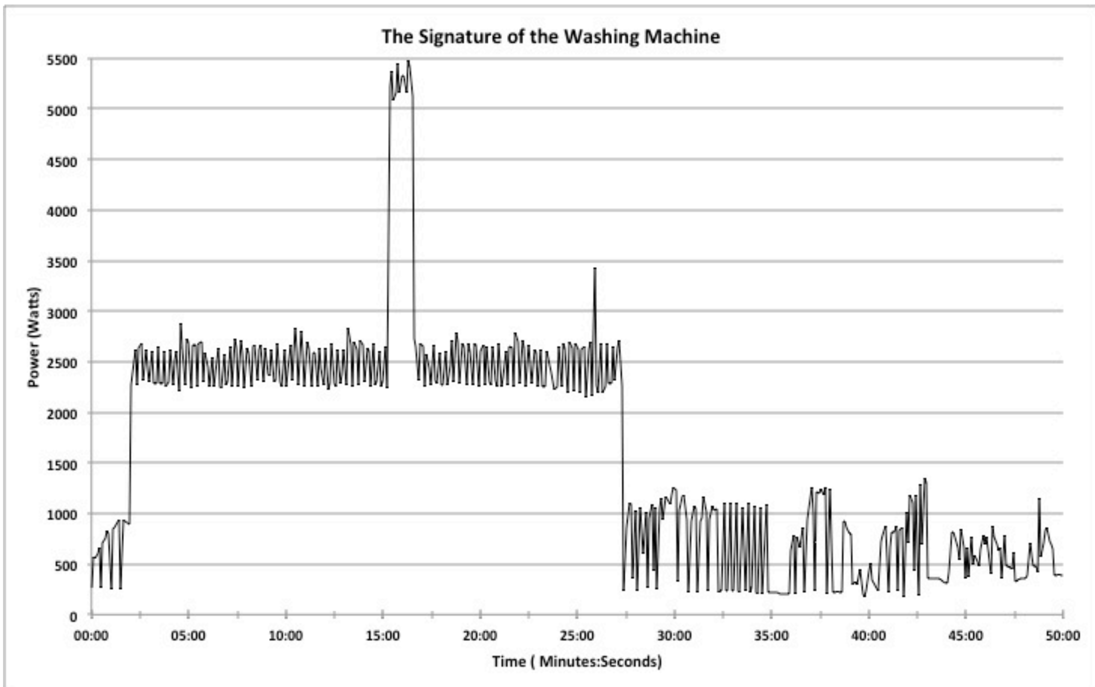

Figure 6.3: A further example of the washing machine signature in household three

In figure 6.3 it is worth noting that the rise in power from 2500 Watts to 5500 Watts at 15 minutes is due to another appliance being turned on (the kettle). As the electricity consumption data for this household is collected using a whole household electricity monitor it is not possible to show only the signature of the washing machine and the two appliances are superimposed on each other. As the model is trained to recognise when the appliance first comes on, the turning on of other appliances when the appliance is on does not affect its recognition providing that two appliances are not turned on with in the same 6-second window (as described in section 6.5.3.1).

As shown by the different signatures for the washing machine in household three in figures 6.2 and 6.3, the power of the washing machine when it first came on was similar, although the overall signatures are quite different, even when taking account of the kettle in figure 6.3.

For appliance recognition, the variability due to the choice of the user is a potential problem when developing a model; however, for the washing machines in household number three (figures 6.3) and the trial household (figures 5.24 and 5.25) both showed different signatures, although when the appliance were first turned on they had similar feature sets. This allowed the washing machine to be trained as one appliance, even though there was evidence of different signatures. This is different from the case of the toaster, section 6.5.6.1, which had to be split into different appliances due to the differences in the signatures when different settings (two-slice and four-slice) were used.

The variability within the signatures of the different washing machines did not affect the ability to train the model for these appliances. For the washing machine, the choices of different modes/cycles and loads are almost endless and the examples given above are for two different instances from two different washing machines in two different households. There is always the possibility that the user could use a setting that produced a different signature when the appliance was first turned on. This would mean that the model would be unable to recognise this instance of the appliance usage. In addition, other washing machines could exhibit very different signatures during different cycles meaning that this assumption would not be valid for those machines. This would also make it much harder to be able to recognise the washing machine, because the model would then have to be trained for a potentially infinite number of signature possibilities, rather than for these two households in which the different cycles were trained as just one appliance for each household.

6.5.6.3. Appliance variability- conclusion

This section has discussed two difference types of appliance variability, which was evident in the analysis of the electricity consumption data from the four households. There was evidence of variability with a number of appliances, either having very different signatures, for example, in the case of the toaster from household three or the electric hob from the trial household or having similar signatures, for example, with the washing machines from household number three and the trial household.

There was also evidence that the occupants of the households were generally habitual in their choices of settings of the appliances that they used. As was evident in the trial household, and described in section 5.6.3, the signatures of the appliances were similar in their feature sets, even though there was the possibility for the appliance signatures to be variable. This was also evident in the appliances from the three households included in this chapter because, although some appliances had the ability to be variable, their signatures all had very similar feature sets.

However, this is only an assumption based on the evidence of similar signatures from the appliance usage, because the occupants were asked only to record appliance usage times, not the settings they used as well. For further work, the recording of the settings of the appliance, as well as the usage times of the appliance, would provide more detail about how signatures of the appliance differ with different settings. It would also provide detail about how habitual the occupants are with the settings of appliances they use.

With the collection of more electricity data, as well as the appliance setting data, a better model could be developed for appliance recognition. As with the small amount of data used to train the models, an assumption was made that the electricity consumption data would be similar in terms of appliance feature set for each use of the appliance. This was found not to be the case for some of the appliances, as discussed in the section 6.5.6.1 and 6.5.6.2, and may have caused problems with recognising the different appliances if they had variable signatures.

The variability within the different appliances did cause problems with the recognition of appliances, although the assumption was made for the development of the models in this research that the users would be habitual in their setting usage and therefore it was assumed that the use of each appliance would consistently produce a similar feature set. This might not be the case for other households and other occupants, meaning that it might not always be possible to recognise appliances used due to variability. This is discussed in the following section.

6.5.7. Differences between households

From the results shown in table 6.18, there was no common window size and feature set that gave the best overall results for each of the households. Some of the reasons behind these differences and also the problems associated with creating a generic model for appliance recognition will be discussed in this section.

For this appliance recognition, a different model was created for each of the households following the approached described in section 5.5.4 and created using the training and test data specific to that household. It was decided not to develop a generic model to recognise the appliances across all of the households as the data and appliances across them were all very different. An example of this is shown by the electricity consumption over seven days across two of the households, shown in figures 6.4 and 6.5.

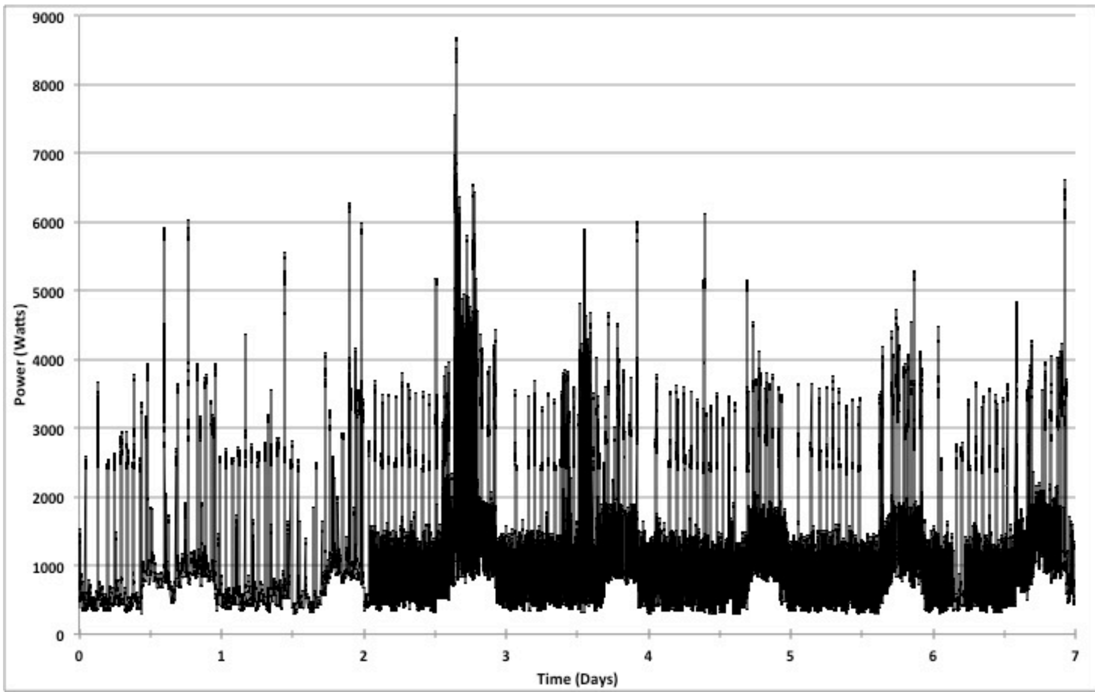

Figure 6.4: 7-day electricity consumption data from one household

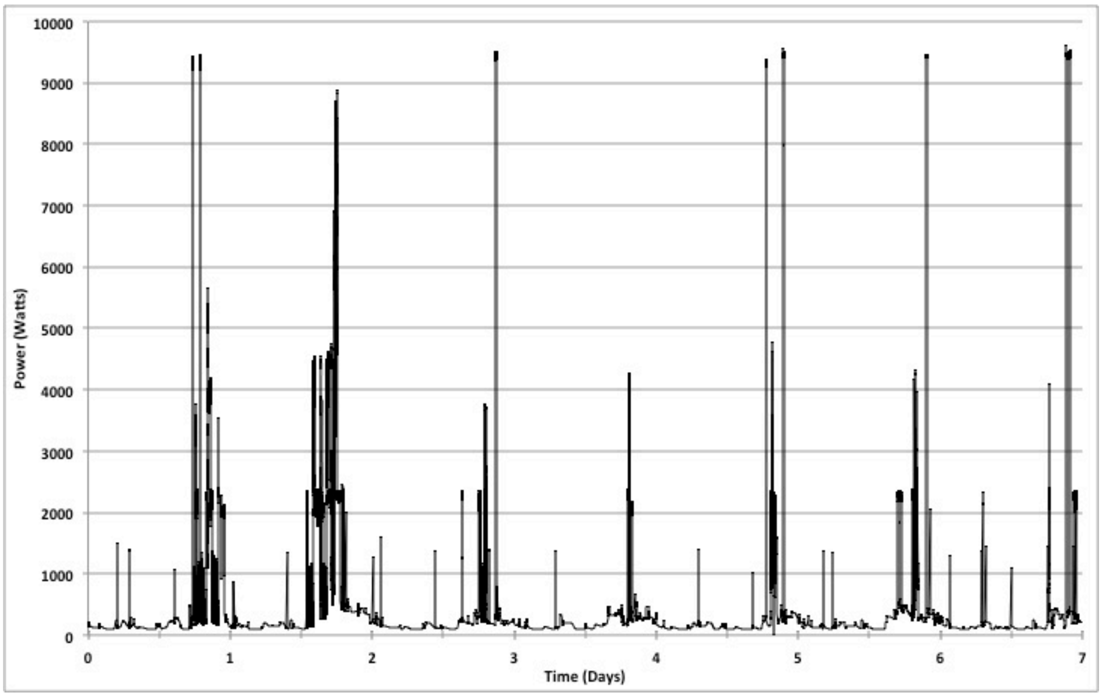

Figure 6.5: 7-day electricity consumption data from another household

As shown by the electricity consumption data in figure 6.4 and 6.5, there were big differences in usage across the households. There are potentially many reasons why the electricity consumption data varies across the households, some examples of these are:

- The size of the household: it is assumed that a larger household would consume more electricity, if only in the background from lighting and background appliances.
- Background appliances, for example, central heating.
- Lights: for this research the use of lights were not recorded as their power is too low.
- Number of occupants: although the number of occupants of the household can increase the overall electricity usage of the household, because this research was focused on recognising appliances and not the overall power consumption of the household the number of occupants was not particularly relevant.

The reasons stated above are separate from this research and will have an effect on the electricity consumption data of any household, to a varying degree. Although there are reasons for the differences in electricity consumption data they would have an effect on developing a single recognition model for appliances across multiple households.

The first of the reasons is due to the differences between appliances of the same type. For this research, each household was asked to record the list of appliances they had in their homes and had used in the data collection period. Although from this recording process there were not that many appliances common to all the households, the kettle being the only appliance that was in use across all the households and was trained and tested to be recognised by the model. The television and washing machine were also shared across the households, although for the reasons discussed in section 6.5.4.1, the television was excluded for all of the households. The washing machine was trained and tested for three of the four households but was excluded from one household for the reasons discussed in section 6.3.1.

While conducting this research, it was found that even though each household might have the same appliance type, e.g. the kettle, the signature of that appliance varied across each of the households, with the lowest power kettle being 2.2 Kilowatts and the highest power kettle being 3 Kilowatts. This is highlighted by the values of the attributes for each of the appliances as shown in the tables in appendix 6. As shown in tables A6.2 and A6.3, the average value for the kettle varies from 1139 Watts for household 1 and 2096 Watts for household 2. This is a relatively large difference in power across one appliance type and the reasons for the differences in power can be put down to different manufacturers. Because of these differences, it was not possible to create a model that would recognise the kettle from all of the households. Therefore, to recognise the appliances from the households a model would have to be created for each of the households, using the training data collected from that household. This makes the recognition of appliances from multiple households more complex and time consuming, as training data would have to be collected from each of the households and a specific model created for each.

Another reason for the differences in the electricity consumption data that affects the ability to recognise appliances is the use of gas for certain appliances. The use of gas by some of the households for general activities, such as cooking, means that this model is missing some activities that could be recognised using other forms of activity recognition, for example IR sensors (Glascock & Kutzik, 2006, 2007).

There is the possibility of extending this work to recognise the use of gas appliances, for example, the cooker, but whole-household gas monitoring does not provide the same granularity of consumption as a whole household electricity monitor due to the differences in how gas appliances function.

To compensate for the use of gas for some of the appliances and recognising certain activities, e.g., cooking, it was the aim to record the usage of some electric appliances that could be used as a proxy for gas appliances. An example for this was the extractor fan being used as a proxy for the use of the gas hob, although as was discussed in section 6.5.4.2, this would require further data. There are also other gas appliances for which a proxy electrical appliance was not available, for example, the gas oven. For households that used gas for their oven there would be a lot of missing data regarding cooking activity. Although though there is the possibility of using a specific gas appliance monitor, as discussed by Intille et al., (2006), this would have detracted from an original aim of this work of being non-intrusive. In addition, the integration of the results from multiple sensors and recording methods would be more complex.

The final reason for the differences in the electricity consumption data that affects the ability to recognise appliances, is the different types of appliances across the households. Although the aim was to record appliances, which were assumed to be common across households, there were only three (two which could be recognised) across the households. More of the appliances were present but used gas, for example, the oven and hob. Due to the differences in the type of appliances within the households, the majority of the households were only trained to recognise four out of the nine appliances from the list.

This has led to the suggestion that further work into monitoring of electricity consumption to infer activities would utilise the monitoring of appliances that the users say they use the most for certain activities, rather than just recording appliances from a generic list that the occupants of the household might not use regularly or which they do not have in their household.

From the discussion of the differences between the electricity consumption data across multiple households it could be argued that the electricity consumption of a household is a unique signature to the household and the occupant(s) of that household. Therefore the method for developing a recognition model for each household would also be specific to that household and the occupants. From the discussion in this section, it is evident that in the case of electricity monitoring, it is not possible to create a generic model to fit all households and appliances.

6.6. Conclusion

This chapter has presented the results from the analysis of the whole house electricity consumption data from three further households. This chapter has also presented a discussion of the results from the analysis conducted in this chapter and the previous analysis conducted in Chapter 5 of this thesis. The next chapter of this thesis (Chapter 7) will provide a conclusion to this thesis.

Chapter 7: Conclusion

7.1. Introduction

This thesis has considered how electricity consumption data might be used to monitor the activities and well-being of people with long-term health conditions and thus help to support them in their own homes. After an initial survey on views of people about monitoring activities in this way, the thesis described the use of a whole house electricity consumption monitor to collect electricity consumption data, which was analysed to provide recognition of the use of specific electrical appliances. The thesis considered how the use of whole house electricity consumption data could be used to support carers and relatives in being able to be aware of the use of appliances by a relative. This chapter will conclude this thesis by drawing together the research undertaken and described in the previous chapters of this thesis. Section 7.2 will discuss the aims and objectives of this research, which were presented in Chapter 1, and outline the extent to which these have been meet. Section 7.3 will discuss the limitations of this research and section 7.4 will highlight the contribution to knowledge from this research. Section 7.5 will discuss areas of further research that have been highlighted by this research and finally section 7.6 will give a final summary of this thesis.

7.2. Achievements of the research aims and objectives

The overall research aim of this thesis, as was highlighted in section 1.6, was to investigate the use of an electricity monitor to recognise and monitor changes in resident's activities. Section 7.2.1 will give an overview of the two main parts of the research and discuss how the research objectives were met. Section 7.2.2 will discuss the key findings from this thesis.

7.2.1. Overview

The first objective of this thesis was "to examine the views of relatives or carers about the use of sensors in the home and specifically activities or tasks that the relative or carer believe are most relevant to be monitored". This objective was met from two parts, the first of these were the literature review (Chapter 2). The literature identified that there was a body of research that discusses the recognising of activities from remote monitoring systems. The literature review also highlighted that there was limited research that had been undertaken on the key activities, and features that carers/relatives would expect from a remote monitoring system intending to aid them in supporting the needs of the person needing support. As these are people who are potential users of the information, it was important to understand their views on what they would like to have information about for reassurance about the elderly/ill person for whom they were a carer or relative.

A survey was therefore conducted to develop a relevant list of activities to monitor, as well as to gain the views of carers and relatives. This survey gathered information about which activities and features the carers and relatives would like to have access to, so as to be reassured about their relative. Chapter four of this thesis presented the analyses of the survey data and provides a discussion of the results. Chapter four thus achieved the first objective of this thesis, as it examined the views of relatives or carers about the use of sensors in the home and specifically the activities or tasks that the relative or carer believed to be most relevant to be monitored. This also contributed to the novelty of the research in this thesis, in this type of survey has not previously been conducted.

The second objective of this thesis was “to examine the feasibility of collecting data from a single electricity monitor within multiple homes”. The literature review identified that there had been previous research with the use of electricity consumption data (collected using various different methods, as highlighted in section 2.4). The literature review had identified previous research that had outlined systems applied to a single home or within a test environment, but no research had investigated the use of whole house electricity consumption monitors in multiple households. Previous research had also not discussed the transferability of recognition models based on a single home or test environment to a larger number of homes. The results from the survey (Chapter 4) were used in this next part to develop a list of electrical appliances, which would be recognised from the whole house electricity consumption data. As described in Chapters 5 and 6, electricity consumption data was collected using a single electricity monitor, from four separate houses, each collecting data for a period of one week. The members of the households were also asked to keep a diary of electrical appliance usage (based on the list of appliances to be monitored, as highlighted in Chapter 3) during this time to use in developing the recognition models. Collecting electricity consumption data from a number of different households addressed the second objective of the thesis, i.e., to examine the feasibility of collecting data from single electricity monitors within different homes. This thesis has provided novel insights into the feasibility of collecting whole house electricity consumption data from this design of single electricity monitor.

The third objective of this thesis was “to analyse the electricity data, from a number of households, to try to determine when different appliances have been used and hence infer different activities have been performed”. To achieve this objective the whole house electricity consumption data and diary data, collected from four houses, was analysed and reported in Chapter 5 and the first part of Chapter 6. The analysis of the electricity consumption data and the discussion of the results in section 6.5, helped to achieve the third objective of this thesis. The analyses of the electricity data from a number of households were used to determine when different appliances had been used and hence to infer different activities that had been performed. This made a novel contribution to research in this area, in that no previous study had collected and analysed electricity data from multiple households.

The final objective of this thesis was “to make recommendations for the use of a single electricity monitor as a remote monitor used to monitor specific activities as well as lifestyle of behavioural changes”. For this objective, the effectiveness of the appliance recognition from each of the houses, as well as the overall observation and issues with the use of whole house electricity consumption data to recognised appliance usage is discussed in section 6.5. Based on this discussion, recommendations about the use of a single electricity monitor as a remote monitor used to monitor specific activities, and hence changes in these activities by someone with a long-term health problem, are highlighted in section 7.4 and 7.5. These recommendations helped to address the final objective of this thesis and make a novel contribution to the literature in this area.

This section gives a summary of the main parts of this thesis and highlights how and to what extend each of the objectives of this thesis were met. The key findings from each part of the thesis are discussed in section 7.2.2.

7.2.2. Key findings

The key findings for this thesis are separated into two sections, which summarise the main findings from the survey (7.2.2.1), and the main findings from the whole house electricity data collection and analyses of these data (7.2.2.2).

7.2.2.1. Survey

The analysis of the responses from the survey produced a number of findings. The first of these findings was that the majority of the respondents classed the list of general activities, provided in the survey, as being either very or quite important to be aware that these activities were being undertaken by their relative. From the survey, the activity of food preparation was identified as the most important general activity and changes in night-time behaviour as the least important. From a list of specific activities that were given in the survey, the majority of the respondents wanted to be made aware that their relative had performed three out of the five specific listed activities, i.e., using the kettle, using the oven and taking medication.

Another key finding from this survey was that the majority of the respondents wanted a remote monitoring system to identify that the person had performed both general and specific activities. The results from the survey also showed that the majority of the respondents thought that it was important that a remote monitoring system was non-intrusive in its monitoring.

The statistical analysis of the survey included undertaking chi-squared tests, which showed that there were very few statistical associations between the caring group, age and gender of participants and the questions answered. This demonstrated that the responses that were provided across groups within the sample were survey were similar.

Another key finding from the survey was from the content analysis of the textual responses from the three open-ended questions. The first question asked the respondents about their three main concerns, which might happen to their relative when they were alone. The results from this analysis showed that falls, the inability to call for help when needed and not eating or drinking properly, were the three most frequently mentioned concerns. The second question asked the respondents which three activities would give them reassurance that their relative was safe and well. The three most frequently occurring responses to this question were regular visits, having contact or a phone call with their relative and monitoring food and drink consumption. The third question asked the respondents for their views on what properties a remote monitoring system should have. The three most frequently occurring responses were timely and reliable alerts or feedback that the system was non-intrusive/unobtrusive/concealable and finally that system was accurate and reliable.

The final key finding from the survey analysis was from the thematic analysis of the first open-ended question. This question asked the respondents of their fears about their relative and a number of themes were identified from the analysis. The major themes from this analysis were related to fears and concerns of their overall health, suffering an accident, their security, their personal well-being and their psychological health.

As highlighted from the literature review in section 2.5, there are two approaches for choosing activities to be monitored using a remote monitoring system; the use of ADLs or choosing a predefined list of activities. As discussed in section 2.5, ADLs were not chosen for this research due to their shortcomings. Instead, the results from this survey were used to inform a list of activities to be monitored based on what activities relatives or carers want to be aware had been undertaken to provide reassurance about the wellbeing of their relative. This is to address the limited research conducted into the views of carers and relatives into what features they may want from a remote monitoring system, as highlighted by the work of Percival & Hanson, (2006). This survey also highlighted requirements of a remote monitoring system, which the development of the appliance recognition model should aim to meet, for example being non-intrusive and providing reliable and accurate feedback. These results then fed into the study of electricity data and helped to develop priorities for the collection and analysis of data.

7.2.2.2. Electricity consumption data collection and analysis

The findings from the data collected on electricity consumption identified that whole house electricity consumption data can be collected from a number of different houses and this data can be effectively transmitted via the Internet to a secure remote storage facility where it can be analysed. The system was easy to install, although the equipment is subject to some limitations within the design and operating conditions. These limitations are highlighted and discussed in more detail in Chapter 3 (section 3.3.2).

The analysis of the electricity consumption data from different houses produced a number of key findings. Perhaps the most important of these findings was that the results from each of the houses varied considerably in terms of their values of positive predictive values (PPV) and sensitivity, with the maximum overall PPV varying between 95% and 54% and the maximum overall sensitivity varying between 90% and 76%. From the previous research discussed in literature review section 2.5, Ruzzelli et al., (2010)'s results had an overall accuracy of 84%. As discussed in section 5.4.6.6, the use of accuracy as a measure of performance was not used for this research due to the large imbalance in the dataset thus this model producing a good value of accuracy (i.e. as shown in table 5.2 with an accuracy for each appliance of >99%) even though not correctly recognising any appliances as being turned on. This makes it difficult to compare the results from this thesis directly with those of Ruzzelli et al., (2010) in any meaningful way. The work of Ruzzelli et al., (2010) also only carried out their experiment from a test environment so could not provide a comparison of the transferability of their method across multiple houses. This thesis therefore contributes to new knowledge by evaluating the results in terms of more meaningful measures, i.e., the PPV and sensitivity, and by undertaking the research in multiple households.

Lines et al., (2011) also provided their results in terms of overall accuracy with a value of 80.32%. As discussed in section 2.4.3.2, data were collected from multiple houses and appliances with the aim of developing a method to automatically recognise appliances rather than in terms of this thesis, which looked at the results based on individual households. Although the results from Lines et al., (2011) gave a good overall accuracy, the results varied between appliances, with some appliances, for example, the oven, washing machine and dishwasher producing poorer results. This was something that was not investigated further by the authors, although it forms part of the discussion and finding of this thesis in terms of transferability between appliances as well as the decision, in this thesis, to use overall PPV and overall sensitivity as a more appropriate measure, rather than overall accuracy.

Lee et al., (2010)'s study gave results for each individual appliance in terms of precision (equivalent to PPV) and recall (equivalent to sensitivity) rather than an overall value. From the tables of results in Chapter 5 and 6, the results for each appliance are provided in terms of PPV and sensitivity although, as discussed in section 5.4.6.6, this research focussed on developing a model to provide a good value of overall PPV and overall sensitivity rather than providing appliance specific values. The work from Lee et al., (2010) was conducted in a test environment and thus could not provide a comparison of the transferability of their method across multiple houses and different appliances.

The data analysis process also showed that the recognition of low power appliances (for example, the television) from the whole house electricity consumption data was not possible using this method of data collection and analysis. From this research only appliances with a power of greater than 800 Watts could be recognised.

Another key finding from this research was the impact of “appliance repeats” as defined in section 5.6.4 of this thesis. These are so-called because the signature of the appliance (i.e. the electric oven and the washing machine), when the user first turns on the appliance, is similar to when the appliance turns itself on and off during its own cycle (as shown in figure 5.5). These “repeats” were an issue with the design of a recognition model as the “repeats” are generally similar to when the appliance is first turned on the model will recognise them as the appliance being turned on again. This initially resulted in a large number of false positives from each of the four houses and a decrease in the PPV results. In terms of recognising activities undertaken by the older/ill person, low PPVs could mean that a carer was informed that a certain activity had taken place, when in fact it had not, thus increasing the risk of a person being left alone when in fact they needed help. This was something that was not highlighted or discussed by Lines et al., (2011) although, as discussed in section 2.5, the result provided in terms of the oven, dishwasher and washing machine gave poorer classification results than other appliances. The finding from this research could be one reason for the poorer classification results for these appliances. However, there could also be other reasons, which are discussed in section 6.5.6, such as appliance variability due to user settings or different appliance signatures due to different manufactures.

To address the problems of repeats, a filter was applied, which remove this “appliance repeats” if there was a previous appliance point, within defined time, the subsequent appliance point would be ignored. The application of this filter produced an improvement in the PPV results. However, the choice of the time of filter is highly subjective and based solely on the repeats of the appliance in question and was different between appliances of the same type, i.e., an electric oven from different households.

The variability of the signature given from the same appliance was another key finding from the analysis of the electricity consumption data. This research showed two types of variability within an appliance. This analysis highlighted the use of appliances, which, depending on the settings chosen by the user, gave a completely different appliance signature depending on that choice. From this data analysis, the use of the toaster gave a completely different signature depending on whether the user was using the two-slice mode or four-slice mode. Although this was just one example in the data, the model had to be trained to recognise the signatures of the toaster for the two-slice mode and the four-slice mode, even though the same appliance was used. Although this is not surprising in hindsight, nonetheless identifying these types of issues is important in developing ways of monitoring activities by analysing electricity consumption. Equally, it is important to be aware that the model can only recognise the signatures of appliances that it has been trained to recognise. Therefore, to recognise the use of appliances and their different signatures accurately, the model would need to be trained carefully and, for some appliances, the variability in their signatures due to different settings would require a large amount of training time. An example of this was given by the attempts at recognition for the electric hob, which could not be recognised due to the variability in its signatures, due to different heat settings (i.e., low to high) collected during this data collection.

The second form of variability was found in signatures which were different in their structure but not in their size when the appliance was first turned on. An example of this was two different signatures from the washing machine in household number three. However, this variability did not affect the ability to train the model to recognise this appliance for the instances that were recorded.

The differences between the electricity consumption data from each of the households were also a key finding of this thesis. Although the differences between the electricity consumption are logical, when considering the variables, such as the house size, the number of occupants or the use of lights, which can affect the electricity usage. This research also highlighted differences between appliances of the same type across the different houses. An example of this is the kettle, the range of power of the kettles across the houses varied from 2.2 Kilowatts to 3 Kilowatts. Although these differences can be explained by there being different models and from different manufacturers, this is also a big range of power across one type of appliance, so models may not be transferable across households.

The differences within the signatures of same appliance across the houses was the reason that an individual model, based on the training data of the appliances from just that house, had to be developed for each of the houses. Although the models do follow the same process in its design, this is more complex and time consuming. The differences between the same appliances, as well as the variability in their signatures, and thus the transferability across different households and appliances, are potentially important issues. This has not been raised previously in any meaningful way by other researchers (as discussed in the literature review section 2.5)

The use of gas has been address as a limitation, as discussed in Chapter 3 of this thesis. The usage of gas appliances varied widely across the houses, with one of the houses using no gas for the appliances to be recognised, whereas other houses used gas for varying appliances from the hob, the oven or the shower. The use of gas appliances limits the appliance usage and information that can be gathered from electricity consumption data and can limit the usefulness of this system.

The final key finding from the analysis of the electricity consumption data was the differences in the types of appliances across the households. Although the list of appliances included several appliances that could be described as common household appliances, only three out of the nine appliances were present in all of the homes as electrical appliances. In the majority of the households, more of the appliances were present but were gas versions of the appliance, for example, the oven. This lead to the majority of the houses having only four out of the nine appliances being able to be trained to be recognised from the electricity consumption data.

This section has highlighted the key findings from the different parts of the research. The next section of this chapter will discuss the limitation of the research.

7.3. Limitations

This section will cover the limitations that have arisen during this research. This section is divided into the limitations that have arisen from the survey and from the electricity consumption data collection and analysis.

7.3.1. Survey analysis

The survey was distributed to a specific ‘research volunteers’ email list at the University of Sheffield. Although sufficient responses were obtained to allow statistical analysis, the response rate could not be determined, because it is not possible to find out how many people are on the list at a given time, as staff join leave the university and can withdraw from the list. However, given the relatively low numbers of responses, out of a staff of over 7,000, it is unlikely that the sample is representative of the University staff population, and certainly not of the wider population of adults in the UK. The generalizability of the results may therefore be limited. Within the survey, almost half (48%) of the respondents of the survey had not cared for an elderly or ill relative. The remaining respondents had either previously cared from an elderly or ill relative (30%) or were currently caring for an elderly or ill relative (22%). The data from the survey were analysed, as highlighted in Chapter 4, based on whether they were at the time, had previously been a carer, or had never been a carer. Within this analysis the number of responses within the groups of those who were currently caring for an elderly or ill relative was relatively small (n= 45) and resulted in some of the chi-squared tests having low expected cell count numbers, meaning that the test results were less reliable. The results had value in providing insights into key issues that concerned people in relation to monitoring systems, and helped focus the design of the second part of the study.

7.3.2. Electricity data collection

The collection of whole house electricity consumption data had a number of limitations, the first of these were with the installation of the equipment. As discussed in more detail in section 3.3.2, to participate in the data collection the participants had to have access to their electricity meter for it to be located within a close distance to their home and have a fixed internet connection in their home. If a participant did not meet any of these requirements, data could not be collected from their home.

The second limitation of the equipment was the monitoring rate. For the electricity monitor placed into the participants' homes (as shown in section 3.3.2) the recording frequency was every 6 seconds. Therefore, this monitor did not provide continuous monitoring of electricity usage and could miss appliances, if they were turned on, and off again, within the 6 seconds, although the use of an appliance in this way is very unlikely. The second limitation of this monitoring frequency was that if two appliances were turned on within the 6-second recording time, it might not be possible to recognise the appliances due to an amalgamation of their signatures. This was highlighted and discussed in more detail in section 5.5.2.1 in relation to the kettle being switched on shortly after another device.

There was also a limitation of the data collection of the appliance usage diary data of incorrect entering of the appliance usage time and date, due to human error. The recording of the microwave highlighted this where the participants entered the wrong date for their usage, in house number one (section 6.2). It is not possible to gauge the extent of this across the households more generally; the study relied on the accuracy of the participants recording the time correctly.

7.3.3. Electricity consumption data analysis

As well as the limitation of the equipment used for the data collection, there were also a number of limitations from the choice of analysis method. The first of these limitations was in the design of the data analysis method, or window design, used for the analysis of the electricity data. For this window design, if two appliances were to be turned on within the same window, the model would not be able to recognise either of the appliances as their feature sets would be distorted and not what the model had been trained to recognise. This is highlighted by the figure 6.1 in section 6.5.3.1 where two appliances are turned or within the same window time. This limitation would therefore have an effect on the overall results of the model, as this point would be classed as a false negative.

The second limitation of the electricity consumption data collection and analysis was a limitation in the design of the training and test datasets. For this research the training and test dataset of appliance usage were based on calendar days rather than the amount of data. This meant that to complete training and testing of an appliance, it had to be used on at least two days of the week. As highlighted in section 6.3.1, the washing machine for household number two had to be excluded, as it was only used on one day of the week. This limitation has lead to a recommendation of further research and this is discussed in section 7.5.

The final limitation highlighted by this data collection and analysis was the limited amount of data relating to the usage of some of the electrical appliances. Some appliances had very low usage, although these did meet the requirements of this method to have a minimum of two training data points and one test data point on at least two days of the week. However, this did lead to some appliances having just one test point with which to test the model. The dates for the training and test datasets had also to be chosen carefully in these cases, so as to be able to provide enough training and test points.

This could lead, as was discussed in section 2.4.5.8, to the model overfitting to the data, because the limited number of data points may not be a representative sample of the data from the wider population (Rogers & Girolami, 2012). However, as described in section 2.4.5.8, cross validation was conducted on the data to address the prospect of potentially overfitting the data. However, the limited data meant that the form of cross validation used was restricted and thus potentially better forms of cross validation (Witten et al., 2011), for example 10-fold cross validation, could not be executed on the data. This is not ideal and lead to a recommendation of further research: this is discussed in section 7.5.

7.4. Significance of the study and contribution to knowledge

This section describes the significance of this study and its contribution to knowledge. This section is organised into two parts, the first part describes the significance and the contribution to knowledge from the survey (7.4.1) and the second part describes the contribution to knowledge and the significance of the electricity consumption data collected and analysis (7.4.2).

7.4.1. Survey of the views of relatives and carers

The literature review in Chapter 2, section 2.5.1, highlighted that there has been only very limited research into the views of carers and relatives into what features and activities they may want monitored from a remote monitoring system. The previous research had consisted of case scenarios used to provide a discussion of the views of carers and relatives, into what features they may want from a remote monitoring system (Percival & Hanson, 2006). The survey carried out as part of this study, as reported in Chapter 4 of this thesis, is therefore the first study of its kind to investigate the views into what activities and features carers and relatives would like to be monitored from a remote monitoring system.

This study found that there are a number of key activities and features that relatives and carers would want from a remote monitoring system. Key amongst these is the need for a remote monitoring system to monitor a range of activities, although as highlighted in the results from the survey, the most important of these activities, is a focus on food and drink, or related activities. The survey also highlighted some important features required by carers and relatives, which had not been previously identified in any study: the first of these was that a remote monitoring system must provide reliable feedback, for example in the form of accurate alerts, so that the number of false alarms from a remote monitoring system are kept to a minimum. The survey also revealed that a remote monitoring system should be non-intrusive, in that the method of recording and collecting the data should not be intrusive to the resident.

The survey reported in Chapter 4 therefore makes a novel contribution in the context of providing a list of activities to be monitored from a remote monitoring system. This is important as it provides a basis for what activities should be monitored by further work in this area. A novel contribution was also provided by highlighting important features that any remote monitoring system should have. This is important as it provides a basis for properties and features that should be included in a remote monitoring system developed in future work in this area.

7.4.2. Electricity data collection and appliance recognition

As highlighted by the literature review in Chapter 2 of this thesis, the data collected from the use of the electricity monitor (as described in Chapter 3) is different from any previous work that has undertaken in this area. For this research, the electricity consumption data were collected using a single whole house electricity monitor and the only electrical variable that was collected by this monitor is the total power consumption of the house, in Watts, collected at a frequency of every 6 seconds. This therefore provides a novel contribution to the understanding of the methods that can be used for capturing electricity consumption data, that can be analysed, collected using a novel non-intrusive method. For this study, these data collections gathered whole house electricity consumption data from different houses, for a period of one week: this has not been undertaken previously. The collection of whole house electricity consumption data (as used for the analysis in Chapters 5 and 6), using this method of collection, from multiple households provides a novel contribution, in that this level of data collection had not been previously undertaken in research. This data collection also provided novel insights into some of the issues that need to be considered when developing methods of appliance recognition, from similar appliances, used in different households, which had not previously been highlighted in earlier research. This therefore provides a methodological contribution through which future theory can be developed for monitoring electricity usage, from multiple households, using a single whole house electricity monitor.

As shown from the results provided in Chapter 5 and Chapter 6 of this thesis, the method developed in this thesis to analyse the electricity consumption data was able to recognise when the occupant had used a range of appliances across the households, although this was with varying levels of accuracy. This is, therefore, a novel contribution to analysing electricity consumption data for recognising appliances. The results from this collection method have not previously been reported in research i.e the use of a single whole house electricity consumption monitor, recording a single electricity variable at a recording interval of 6 seconds.

As shown in Chapter 5, the analysis of the electricity consumption data involved the development of an analysis method, taking forward the concept based on the method by Lee et al., (2010). The forward and backwards sliding window method developed for the construction of a feature set from the whole house electricity consumption data, as described in section 5.5.1, had not previously been used and presents a new approach to analysing whole house electricity consumption data.

As shown in the discussion in section 6.5 and highlighted in the key findings in section 7.2.2.2, there are differences in the signatures of the same appliance type across the different houses. As well as the differences within the same type of appliances, the electricity consumption of a different house varies widely. Therefore, this research has provided a novel contribution by showing that it is not possible to create a generic recognition model for the recognition of appliances across different homes and also argues that the electricity consumption data is potentially unique to the occupants and the households and therefore cannot be compared. As shown by the data analysis in Chapters 5 and 6, individual models for each house are therefore required. This finding is novel, as it has not previously been reported in the literature, as well as providing a basis of a method for recognising appliances across multiple households, using the data collection method undertaken in this thesis.

This thesis is helping make sense of what activities should be monitor from a remote monitoring systems. As well as highlighting how electricity data can be collected and monitored, in a non-intrusive way, so that people's activities can be used to highlight potential changes in health and well-being. This further helps our understanding of how older people and those with long-term health conditions can be better supported to live longer in their own homes.

7.5. Recommendation for further research

Based on this research process and the analysis of the data, a number of areas of further research have been highlighted.

Within this research the use of gas appliances have resulted in a limitation with the method, as information about appliance usage and activities is missing. A proposed area of future work is investigating the combination of electricity consumption monitor with a gas consumption monitor to give a comprehensive overview of appliance usage within a home.

The list of electrical appliance used for this research highlighted the differences between what could be considered as “common household appliances”, because some of the appliances were not present, or the occupant rarely used them or they were gas appliances. From this research it is recommended that the choice of appliances to be monitored should be based on what appliances the occupant uses most, for each activity, instead of a generic list of activities.

Due to the low appliance usage that was highlighted during this thesis, it is recommended that the length of future data collection periods be increased. This would allow for the capturing of more data, which could be used to train and test the models further. A longer collection time could also be used to investigate the differences within appliance signatures (as highlighted in section 6.5) and their effect on the overall recognition of that appliance.

Collecting longitudinal data could also be used to highlight changes in activities, which could be attributed to changes in health. An example of this is given by figure 7.1. Figure 7.1 shows the plot of recognised appliance usage, based on the results from the model developed for the trial house, as described in section 5.5.5. This plot shows three days of appliance usage, based on the model’s classifications and could be used to show trends in appliance usage patterns and which may be used to show changes in these patterns, if more data were available. Already, from three days of data, it is possible to show that the resident uses their microwave at between 6am and 7am everyday, that they had not used any of the recognised appliances, between the hours of 7am and 5pm, so it could be inferred that they were not in the house between these hours and that they did not use any recognised appliances after the hours of 10pm, so it might be inferred that they go to sleep at that time. This information could then be used to identify patterns of behaviour, and therefore, to identify when such patterns were interrupted among people living on their own, which might indicate that they were in need of help.

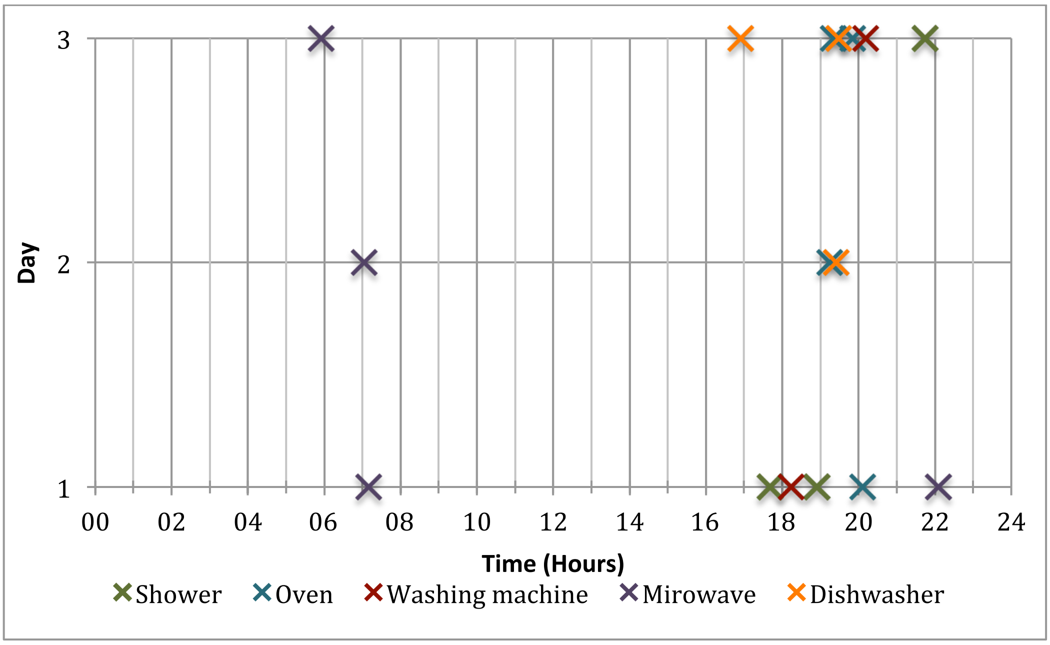

Figure 7.1: Plot of appliance usage for three days

7.6. Summary

This research has highlighted that it is possible to use a single whole house electricity monitor, to show the performing of activities in a non-invasive way. It was also highlighted that there are large differences between appliances and electricity usage across different households, which made the process of recognising appliances from each of the households an individual process and, in some cases, very time consuming.

More research is needed to determine whether the use of just an electricity monitor can be used to show enough detail in the performing of activities or if an electricity monitor needs to be combined with other types of monitoring sensors. The collection and analysis of longitudinal whole house electricity consumption data is also another area, which needs further investigation.

Privacy, in terms of electricity usage monitoring, is an area that the researcher feels needs to be highlighted and addressed more in the future. This is especially the case, now that the wealth of information, which can be gathered from a household's electricity usage, more needs to be done to highlight how this information should and should not be used and to protect the privacy etc., of those from whom the data were obtained.

References

Adlam, T., Faulkner, R., Orpwood, R., Jones, K., Macijauskiene, J., & Budraitiene, A. (2004). The installation and support of internationally distributed equipment for people with dementia. *IEEE Transactions on Information Technology in Biomedicine*, *8*(3), 253–257. http://doi.org/10.1109/TITB.2004.834393

Adriaans, P., & Zantinge, D. (1996). *Data mining*. Harlow: Addison-Wesley.

Agoulmine, N., Deen, M. J., Lee, J.-S., & Meyyappan, M. (2011). U-Health Smart Home. *IEEE Nanotechnology Magazine*, *5*(3), 6–11. http://doi.org/10.1109/MNANO.2011.941951

Al Ameen, M., Liu, J., & Kwak, K. (2012). Security and privacy issues in wireless sensor networks for healthcare applications. *Journal of Medical Systems*, *36*(1), 93–101. http://doi.org/10.1007/s10916-010-9449-4

Alam, M. R., Reaz, M. B. I., & Ali, M. A. M. (2012). A Review of Smart Homes-Past, Present, and Future. *IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews*, *42*(6), 1190–1203. http://doi.org/10.1109/TSMCC.2012.2189204

Alpaydin, E. (2010). *Introduction to Machine Learning* (2nd ed.). Cambridge, MA; London, England: MIT Press. Retrieved from http://ieeexplore.ieee.org/xpl/bkabstractplus.jsp?bkn=6267367

Altman, D. (1999). *Practical statistics for medical research* (1st ed.). London ; New York: Chapman and Hall.

Alwan, M., Dalal, S., Mack, D., Kell, S. W., Turner, B., Leachtenauer, J., & Felder, R. (2006). Impact of monitoring technology in assisted living: outcome pilot. *IEEE Transactions on Information Technology in Biomedicine : A Publication of the IEEE Engineering in Medicine and Biology Society*, *10*(1), 192–198. http://doi.org/10.1109/TITB.2005.855552

Anliker, U., Ward, J. A., Lukowicz, P., Tröster, G., Dolveck, F., Baer, M., … Vuskovic, M. (2004). AMON: a wearable multiparameter medical monitoring and alert system. *IEEE Transactions on Information Technology in Biomedicine : A Publication of the IEEE Engineering in Medicine and Biology Society*, *8*(4), 415–27. http://doi.org/10.1109/TITB.2004.837888

Apté, C. (1997). Data mining: an industrial research perspective. *IEEE Computational Science & Engineering*, *4*(2), 6–9. http://doi.org/10.1109/99.609825

Apté, C., & Weiss, S. (1997). Data mining with decision trees and decision rules. *Future Generation Computer Systems*, *13*(2–3), 197–210. http://doi.org/10.1016/S0167-739X(97)00021-6

Barlow, J., Singh, D., Bayer, S., & Curry, R. (2007). A systematic review of the benefits of home telecare for frail elderly people and those with long-term conditions. *Journal of Telemedicine and Telecare*, *13*(4), 172–179. http://doi.org/10.1258/135763307780908058

Barron, L. (2006). Paradigms. In V. Jupp (Ed.), *The SAGE Dictionary of Social Research Methods* (pp. 212–214). London, England; Thousand Oaks, California: SAGE Publications.

Batista, G. E. A. P. A., Prati, R. C., & Monard, M. C. (2004). A Study of the Behavior of Several Methods for Balancing Machine Learning Training Data. *SIGKDD Explor. Newsl.*, *6*(1), 20–29. http://doi.org/10.1145/1007730.1007735

Belkin International Inc. (2012). Conserve Insight. Retrieved January 2015 from http://www.belkin.com/conserve/insight/

Belley, C., Gaboury, S., Bouchard, B., & Bouzouane, A. (2013). Activity recognition in smart homes based on electrical devices identification. In *Proceedings of the 6th International Conference on PErvasive Technologies Related to Assistive Environments - PETRA '13* (pp. 1–8). New York: ACM Press. http://doi.org/10.1145/2504335.2504342

Belley, C., Gaboury, S., Bouchard, B., & Bouzouane, A. (2014). An efficient and inexpensive method for activity recognition within a smart home based on load signatures of appliances. *Pervasive and Mobile Computing*, *12*, 58–78. http://doi.org/10.1016/j.pmcj.2013.02.002

Berenguer, M., Giordani, M., Giraud-By, F., & Noury, N. (2008). Automatic detection of activities of daily living from detecting and classifying electrical events on the residential power line. In *10th International Conference on e-health Networking, Applications and Services, 2008. HealthCom 2008.* (pp. 29–32). IEEE. http://doi.org/10.1109/HEALTH.2008.4600104

Blaxter, L., Hughes, C., & Tight, M. (2010). *How to research* (4th ed.). Maidenhead: McGraw-Hill/Open University Press. Retrieved from https://www.dawsonera.com/abstract/9780335238699

Botsis, T., & Hartvigsen, G. (2008). Current status and future perspectives in telecare for elderly people suffering from chronic diseases. *Journal of Telemedicine and Telecare*, *14*(4), 195–203. http://doi.org/10.1258/jtt.2008.070905

Bowes, A., Dawson, A., & Bell, D. (2012). Ethical Implications of Lifestyle Monitoring Data in Ageing Research. *Information, Communication & Society*, *15*(1), 5–22. http://doi.org/10.1080/1369118X.2010.530673

Brownsell, S., Blackburn, S., & Hawley, M. S. (2008). An evaluation of second and third generation telecare services in older people's housing. *Journal of Telemedicine and Telecare*, *14*(1), 8–12. http://doi.org/10.1258/jtt.2007.070410

Brownsell, S., Bradley, D., Blackburn, S., Cardinaux, F., & Hawley, M. S. (2011). A systematic review of lifestyle monitoring technologies. *Journal of Telemedicine and Telecare*, *17*(4), 185–189. http://doi.org/10.1258/jtt.2010.100803

Bryman, A. (1988). *Quantity and quality in social research*. London: Unwin Hyman.

Bryman, A. (2012). *Social Research Methods* (4th ed.). Oxford ; New York: Oxford University Press.

Bryman, A., & Cramer, D. (2011). *Quantitative Data Analysis with IBM SPSS 17, 18 & 19: A Guide for Social Scientists*. London: Routledge.

Bucks, R. S., Ashworth, D. L., Wilcock, G. K., & Siegfried, K. (1996). Assessment of activities of daily living in dementia: development of the Bristol Activities of Daily Living Scale. *Age and Ageing*, *25*(2), 113–120. http://doi.org/10.1093/ageing/25.2.113

Chan, H., & Perrig, A. (2003). Security and Privacy in Sensor Networks. *Computer*, *36*(10), 103–105. http://doi.org/10.1109/MC.2003.1236475

Chan, M., Campo, E., Estève, D., & Fourniols, J.-Y. (2009). Smart homes - current features and future perspectives. *Maturitas*, *64*(2), 90–7. http://doi.org/10.1016/j.maturitas.2009.07.014

Chan, M., Estève, D., Escriba, C., & Campo, E. (2008). A review of smart homes- present state and future challenges. *Computer Methods and Programs in Biomedicine*, *91*(1), 55–81. http://doi.org/10.1016/j.cmpb.2008.02.001

Chen, L., Lu, S., & Ram, J. (2004). Compressed pattern matching in DNA sequences. In *Proceedings of 2004 IEEE Computational Systems Bioinformatics Conference, 2004. CSB 2004.* (pp. 62–68). IEEE. http://doi.org/10.1109/CSB.2004.1332418

Chetty, M., Tran, D., & Grinter, R. E. (2008). Getting to Green: Understanding Resource Consumption in the Home. In *Proceedings of the 10th International Conference on Ubiquitous Computing - UbiComp '08* (pp. 242–251). New York: ACM Press. http://doi.org/10.1145/1409635.1409668

Climate Change Act 2008. (c.27). London: The Stationery Office. Retrieved 4/3/15 from http://www.legislation.gov.uk/ukpga/2008/27/pdfs/ukpga_20080027_en.pdf

Cracknell, R. (2010). *The Ageing Population. Key Issues for the New Parliament 2010*. Retrieved from http://www.parliament.uk/documents/commons/lib/research/key_issues/Key-Issues-The-ageing-population2007.pdf

Creswell, J. (2014). *Research design : Qualitative, quantitative, and mixed method approaches* (4th ed.). Thousand Oaks: SAGE Publications.

Cristianini, N., & Shawe-Taylor, J. (2000). *An introduction to Support Vector Machines : and other kernel-based learning methods*. Cambridge: Cambridge University Press.

Current Cost. (2015). EnviR. Retrieved January 2015 from http://www.currentcost.com/product-envir.html

Dang, N. Le, Le, D., & Le, V. T. (2016). A New Multiple-Pattern Matching Algorithm for the Network Intrusion Detection System. *International Journal of Engineering and Technology*, *8*(2), 94–100. http://doi.org/10.7763/IJET.2016.V8.865

Das, S., Cook, D., Battacharya, A., Heierman III, E. O., & Lin, T. Y. (2002). The Role of Prediction algorithms in the Mavhome Smart Home Architecture. *IEEE Wireless Communications*, *9*(6), 77–84. http://doi.org/10.1109/MWC.2002.1160085

De Ruyter, B., & Aarts, E. (2004). Ambient intelligence. In *Proceedings of the working conference on Advanced visual interfaces - AVI '04* (p. 203). New York: ACM Press. http://doi.org/10.1145/989863.989897

Demiris, G., & Hensel, B. (2008). Technologies for an aging society: a systematic review of "smart home" applications. *IMIA Yearbook 2008: Access to Health Information*, *3*(1), 33–40. Retrieved from http://www.schattauer.de/t3page/1214.html?manuscript=9816&L=1

Demiris, G., Hensel, B. K., Skubic, M., & Rantz, M. (2008). Senior residents' perceived need of and preferences for "smart home" sensor technologies. *International Journal of Technology Assessment in Health Care*, *24*(1), 120–4. http://doi.org/10.1017/S0266462307080154

Demiris, G., Rantz, M., Aud, M., Marek, K., Tyrer, H., Skubic, M., & Hussam, A. (2004). Older adults' attitudes towards and perceptions of "smart home" technologies: a pilot study. *Medical Informatics and the Internet in Medicine*, *29*(2), 87–94. http://doi.org/10.1080/14639230410001684387

Demongeot, J., Virone, G., Duchêne, F., Benchetrit, G., Hervé, T., Noury, N., & Rialle, V. (2002). Multi-sensors acquisition, data fusion, knowledge mining and alarm triggering in health smart homes for elderly people. *Comptes Rendus Biologies*, *325*(6), 673–682. http://doi.org/10.1016/S1631-0691(02)01480-4

Draper, H., & Sorell, T. (2013). Telecare, remote monitoring and care. *Bioethics*, *27*(7), 365–72. http://doi.org/10.1111/j.1467-8519.2012.01961.x

Duda, R., Hart, P. E., & Stork, D. G. (2001). *Pattern classification* (2nd ed.). New York: Wiley.

Ebling, J., & Scheuermann, G. (2003). Clifford Convolution and Pattern Matching on Vector Fields. In *Proceedings of the IEEE Visualization Conference* (pp. 193–200). IEEE. http://doi.org/10.1109/VISUAL.2003.1250372

Efergy Technologies Limited. (2014a). Energy monitoring socket 2.0. Retrieved January 2015 from http://efergy.com/eu/energy-monitoring-socket

Efergy Technologies Limited. (2014b). e2 classic. Retrieved January 2015 from http://efergy.com/eu/e2v2-monitor

Energy Saving Trust. (2012). Powering the Nation: Household electricity-using habits revealed. Retrieved from http://www.energysavingtrust.org.uk/sites/default/files/reports/PoweringthenationreportCO332.pdf

Farinaccio, L., & Zmeureanu, R. (1999). Using a pattern recognition approach to disaggregate the total electricity consumption in a house into the major end-uses. *Energy and Buildings*, *30*, 245–259. http://doi.org/10.1016/S0378-

7788(99)00007-9

Fayyad, U., Piatetsky-Shapiro, G., & Smyth, P. (1996a). From data mining to knowledge discovery in databases. *AI Magazine*, *17*(3), 37–54. http://doi.org/http://dx.doi.org/10.1609/aimag.v17i3.1230

Fayyad, U., Piatetsky-Shapiro, G., & Smyth, P. (1996b). Knowledge discovery and data mining: towards a unifying framework. In *2nd International Conference on Knowledge Discovery and Data Mining* (pp. 82–88). Retrieved from https://www.aaai.org/Papers/KDD/1996/KDD96-014.pdf

Fischer, C. (2008). Feedback on household electricity consumption: a tool for saving energy? *Energy Efficiency*, *1*(1), 79–104. http://doi.org/10.1007/s12053-008-9009-7

Fleury, A., Vacher, M., & Noury, N. (2010). SVM-based multimodal classification of activities of daily living in Health Smart Homes: sensors, algorithms, and first experimental results. *IEEE Transactions on Information Technology in Biomedicine*, *14*(2), 274–283. http://doi.org/10.1109/TITB.2009.2037317

Flick, U. (2014). *An introduction to qualitative research* (5th ed.). Los Angeles: SAGE Publications.

Flick, U., Von Kardorff, E., & Steinke, I. (2004). What is qualitative research? An introduction to the field. In U. Flick, E. Von Kardorff, & I. Steinke (Eds.), *A Companion to Qualitative Research* (pp. 3–11). London;Thousand Oaks, Calif.: SAGE Publications.

Foster, L., Diamond, I., & Jefferies, J. (2015). *Beginning statistics: An introduction for social scientists* (2nd ed.). Los Angeles: SAGE Publications.

Franco, G. C., Gallay, F., Berenguer, M., Mourrain, C., & Couturier, P. (2008). Non-invasive monitoring of the activities of daily living of elderly people at home - a pilot study of the usage of domestic appliances. *Journal of Telemedicine and Telecare*, *14*(5), 231–5. http://doi.org/10.1258/jtt.2008.071207

García, S., Fernández, A., & Herrera, F. (2009). Enhancing the effectiveness and interpretability of decision tree and rule induction classifiers with evolutionary training set selection over imbalanced problems. *Applied Soft Computing Journal*, *9*(4), 1304–1314. http://doi.org/10.1016/j.asoc.2009.04.004

Ghahramani, Z., & Jordan, M. (1997). Factorial hidden Markov models. *Machine Learning*, *29*(2), 245–273. http://doi.org/10.1023/A:1007425814087

Glascock, A., & Kutzik, D. (2006). The Impact of Behavioral Monitoring Technology on the Provision of Health Care in the Home. *Journal of Universal Computer Science*, *12*(1), 59–79. http://doi.org/10.3217/jucs-012-01-0059

Glascock, A., & Kutzik, D. (2007). An Evidentiary Study of the Uses of Automated Behavioral Monitoring. In *21st International Conference on Advanced Information Networking and Applications Workshops, 2007, AINAW '07.* (Vol. 2, pp. 858–862). IEEE. http://doi.org/10.1109/AINAW.2007.81

Green Energy Options Ltd. (n.d.) Retrieved January 2015 from http://www.greenenergyoptions.co.uk/support/ensemble/

Gupta, S., Reynolds, M. S., & Patel, S. N. (2010). ElectriSense : Single-Point Sensing Using EMI for Electrical Event Detection and Classification in the Home. In *Proceedings of the 12th ACM International Conference on Ubiquitous Computing* (pp. 139–148). New York: ACM. http://doi.org/10.1145/1864349.1864375

Gurney, K. (1997). *An introduction to neural networks*. London: UCL Press.

Hagras, H., Callaghan, V., Colley, M., Clarke, G., Pounds-Cornish, A., & Duman, H. (2004). Creating an ambient-intelligence environment using embedded agents. *IEEE Intelligent Systems*, *19*(6), 12–20. http://doi.org/10.1109/MIS.2004.61

Hand, D., Mannila, H., & Smyth, P. (2001). *Principles of data mining*. Cambridge, Mass.: MIT Press. Retrieved from http://ieeexplore.ieee.org/xpl/bkabstractplus.jsp?bkn=6267275

Hanson, J., Osipovič, D., Hinew, N., Amaral, T., Curry, R., & Barlow, J. (2007). Lifestyle monitoring as a predictive tool in telecare. *Journal of Telemedicine and Telecare*, *13*(suppl 1), 26–28. http://doi.org/10.1258/135763307781645040

Hart, G. W. (1992). Nonintrusive appliance load monitoring. *Proceedings of the IEEE*, *80*(12), 1870–1891. http://doi.org/10.1109/5.192069

Haykin, S. (1999). *Neural networks : A comprehensive foundation*. (2nd ed., Ed.). Saddle River, N.J. ; London: Prentice Hall International.

Helal, A., Cook, D. J., & Schmalz, M. (2009). Smart Home-Based Health Platform for Behavioral Monitoring and Alteration of Diabetes Patients. *Journal of Diabetes Science and Technology*, *3*(1), 141–148. http://doi.org/10.1177/193229680900300115

Helal, S., Mann, W., El-Zabadani, H., King, J., Kaddoura, Y., & Jansen, E. (2005). The gator tech smart house: A programmable pervasive space. *Computer*, *38*(3), 50–60. http://doi.org/10.1109/MC.2005.107

Howell, K. (2013). *An Introduction to the Philosophy of Methodology*. London: SAGE Publications. Retrieved from http://srmo.sagepub.com/view/an-introduction-to-the-philosophy-of-methodology/SAGE.xml

Intille, S., Larson, K., Tapia, E., Beaudin, J., Kaushik, P., Nawyn, J., & Rockinson, R. (2006). Using a Live-In Laboratory for Ubiquitous Computing Research. In K. Fishkin, B. Schiele, P. Nixon, & A. Quigley (Eds.), *Pervasive Computing* (Vol. 3968, pp. 349–365). Springer Berlin Heidelberg. http://doi.org/10.1007/11748625_22

Jain, A. K. (2010). Data clustering: 50 years beyond K-means. *Pattern Recognition Letters*, *31*(8), 651–666. http://doi.org/10.1016/j.patrec.2009.09.011

Jain, A. K., Duin, R., & Mao, J. (2000). Statistical pattern recognition: A review. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *22*(1), 4–37.

http://doi.org/10.1109/34.824819

Kantardzic, M. (2011). *Data mining: concepts, models, methods, and algorithms* (2nd ed.). Hoboken, NJ: Wiley-IEEE Press. Retrieved from http://ieeexplore.ieee.org/xpl/bkabstractplus.jsp?bkn=6105606

Kargl, F., Lawrence, E., Fischer, M., & Lim, Y. Y. (2008). Security, Privacy and Legal Issues in Pervasive eHealth Monitoring Systems. In *7th International Conference on Mobile Business* (pp. 296–304). IEEE. http://doi.org/10.1109/ICMB.2008.31

Katz, S., Downs, T. D., Cash, H. R., & Grotz, R. C. (1970). Progress in development of the index of ADL. *The Gerontologist*, *10*(1), 20–30. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/5420677

Kawarada, A., Takagi, T., Tsukada, A., Sasaki, K., Ishijima, M., Tamura, T., … Yamakoshi, K. (1998). Evaluation of automated health monitoring system at the "Welfare Techno House." In *Proceedings of the 20th Annual International Conference of the IEEE Engineering in Medicine and Biology Society* (Vol. 20, pp. 1984–1987). IEEE. http://doi.org/10.1109/IEMBS.1998.746992

Kidd, C., Orr, R., Abowd, G., Atkeson, C., Essa, I., MacIntyre, B., … Newstetter, W. (1999). The aware home: A living laboratory for ubiquitous computing research. In *Proceedings of the Second International Workshop on Cooperative Buildings, Integrating Information, Organization, and Architecture* (pp. 191–198). Springer. Retrieved from http://dl.acm.org/citation.cfm?id=645969.674887

Kim, J., Choi, H., Wang, H., Agoulmine, N., Deerv, M. J., & Hong, J. W. (2010). POSTECH's U-Health Smart Home for elderly monitoring and support. In *2010 IEEE International Symposium on "A World of Wireless, Mobile and Multimedia Networks" (WoWMoM)* (pp. 1–6). IEEE. http://doi.org/10.1109/WOWMOM.2010.5534977

Kim, Y., Schmid, T., Charbiwala, Z. M., & Srivastava, M. B. (2009). ViridiScope: design and implementation of a fine grained power monitoring system for homes. In *Proceedings of the 11th International Conference on Ubiquitous Computing* (pp. 245–254). ACM. http://doi.org/10.1145/1620545.1620582

Kim, Y., Schmid, T., Srivastava, M. B., & Wang, Y. (2009). Challenges in resource monitoring for residential spaces. In *Proceedings of the First ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings* (pp. 1–6). ACM. http://doi.org/10.1145/1810279.1810281

Koch, S. (2006). Home telehealth- current state and future trends. *International Journal of Medical Informatics*, *75*(8), 565–76. http://doi.org/10.1016/j.ijmedinf.2005.09.002

Kolter, J. Z., & Johnson, M. J. (2011). REDD : A Public Data Set for Energy Disaggregation Research. In *Workshop on Data Mining Applications in Sustainability (SIGKDD)* (Vol. 20, pp. 59–62). ACM. Retrieved from http://redd.csail.mit.edu/kolter-kddsust11.pdf

Kotsiantis, S. B. (2013). Decision trees: A recent overview. *Artificial Intelligence Review*, *39*(4), 261–283. http://doi.org/10.1007/s10462-011-9272-4

Kotsiantis, S. B., Zaharakis, I. D., & Pintelas, P. E. (2006). Machine learning: A review of classification and combining techniques. *Artificial Intelligence Review*, *26*(3), 159–190. http://doi.org/10.1007/s10462-007-9052-3

Kröse, B., van Kasteren, T., Gibson, C., & van den Dool, T. (2008). Care: Context awareness in residences for elderly. In *The 6th International Conference of the International Society for Gerontechnology* (pp. 101–105). Retrieved from http://dare.uva.nl/record/289919

Krumm, J., Harris, S., Meyers, B., Brumitt, B., Hale, M., & Shafer, S. (2000). Multi-camera multi-person tracking for easyliving. In *Proceedings of the Third IEEE International Workshop on Visual Surveillance, 2000* (pp. 3–10). IEEE. http://doi.org/10.1109/VS.2000.856852

Kulkarni, A. S., Welch, K. C., & Harnett, C. K. (2011). A review of electricity monitoring and feedback systems. In *Proceedings of IEEE Southeastcon, 2011* (Vol. 1, pp. 321–326). IEEE. http://doi.org/10.1109/SECON.2011.5752958

Kurgan, L. A., & Musilek, P. (2006). A survey of Knowledge Discovery and Data Mining process models. *The Knowledge Engineering Review*, *21*(1), 1–24. http://doi.org/10.1017/S0269888906000737

Landau, R., Werner, S., Auslander, G. K., Shoval, N., & Heinik, J. (2009). Attitudes of Family and Professional Care-Givers towards the Use of GPS for Tracking Patients with Dementia: An Exploratory Study. *British Journal of Social Work*, *39*(4), 670–692. http://doi.org/10.1093/bjsw/bcp037

Lapadat, J. C. (2010). Thematic Analysis. In A. Mills, G. Durepos, & E. Wiebe (Eds.), *Encyclopedia of Case Study Research* (pp. 926–928). Thousand Oaks, CA: SAGE Publications. http://doi.org/http://dx.doi.org/10.4135/9781412957397

Larose, D. (2005). *Discovering knowledge in data : an introduction to data mining*. Hoboken, N.J.: Wiley-Interscience.

Lawton, M. P., & Brody, E. M. (1969). Assessment of older people: self-maintaining and instrumental activities of daily living. *The Gerontologist*, *9*(3), 179–86. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/5349366

LeBellego, G., Noury, N., Virone, G., Mousseau, M., & Demongeot, J. (2006). A model for the measurement of patient activity in a hospital suite. *IEEE Transactions on Information Technology in Biomedicine*, *10*(1), 92–9. http://doi.org/10.1109/TITB.2005.856855

Lee, P. (2004). *Bayesian statistics : an introduction* (3rd ed.). London: Hodder Arnold.

Lee, S., Lin, G., Jih, W., & Hsu, J. Y. (2010). Appliance Recognition and Unattended Appliance Detection for Energy Conservation. In *2010 AAAI Workshop on Plan, Activity, and Intent Recognition (PAIR)* (pp. 37–44). Retrieved from http://www.aaai.org/ocs/index.php/WS/AAAIW10/paper/view/2012

Li Jiang, Da-You Liu, & Bo Yang. (2004). Smart home research. In *Proceedings of 2004 International Conference on Machine Learning and Cybernetics* (Vol. 2, pp. 659–663). IEEE. http://doi.org/10.1109/ICMLC.2004.1382266

Lin, G., Lee, S., Hsu, J. Y., & Jih, W. (2010). Applying Power Meters for Appliance Recognition on the Electric Panel. In *the 5th IEEE Conference on Industrial Electronics and Applications (ICIEA), 2010* (pp. 2254–2259). IEEE. http://doi.org/10.1109/ICIEA.2010.5515385

Lines, J., Bagnall, A., Caiger-Smith, P., & Anderson, S. (2011). Classification of Household Devices by Electricity Usage Profiles. In *Intelligent Data Engineering and Automated Learning - IDEAL 2011* (Vol. 6936 LNCS, pp. 403–412). Berlin, Heidelberg: Springer Berlin Heidelberg. http://doi.org/10.1007/978-3-642-23878-9_48

Liu, W., Chawla, S., Cieslak, D., & Chawla, N. V. (2010). A Robust Decision Tree Algorithm for Imbalanced Data Sets. In *Proceedings of the Tenth SIAM International Conference on Data Mining* (pp. 766–777). Retrieved from http://epubs.siam.org/doi/pdf/10.1137/1.9781611972801.67

Lotfi, A., Langensiepen, C., Mahmoud, S. M., & Akhlaghinia, M. J. (2011). Smart homes for the elderly dementia sufferers: identification and prediction of abnormal behaviour. *Journal of Ambient Intelligence and Humanized Computing*, *3*(3), 205–218. http://doi.org/10.1007/s12652-010-0043-x

Magnusson, L., & Hanson, E. J. (2003). Ethical issues arising from a research , technology and development project to support frail older people and their family carers at home. *Health & Social Care in the Community*, *11*(5), 431–439. http://doi.org/10.1046/j.1365-2524.2003.00446.x

Marsland, S. (2009). *Machine learning : an algorithmic perspective*. Boca Raton: CRC Press.

Mazurowski, M. a., Habas, P. a., Zurada, J. M., Lo, J. Y., Baker, J. a., & Tourassi, G. D. (2008). Training neural network classifiers for medical decision making: The effects of imbalanced datasets on classification performance. *Neural Networks*, *21*(2-3), 427–436. http://doi.org/10.1016/j.neunet.2007.12.031

McNeill, P., & Chapman, S. (2005). *Research methods* (3rd ed.). London; New York: Routledge. Retrieved from https://www.dawsonera.com/abstract/9780203463000

McQueen, R. A., & Knussen, C. (2002). *Research methods for social science: A practical introduction*. Harlow: Prentice Hall.

Meingast, M., Roosta, T., & Sastry, S. (2006). Security and Privacy Issues with Health Care Information Technology. In *28th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, 2006* (pp. 5453–5458). IEEE. http://doi.org/10.1109/IEMBS.2006.260060

Milligan, C., Roberts, C., & Mort, M. (2011). Telecare and older people: who cares where? *Social Science & Medicine*, *72*(3), 347–54. http://doi.org/10.1016/j.socscimed.2010.08.014

Mitchell, T. (1997). *Machine learning*. New York ; London: McGraw-Hill.

Mozer, M. (1998). The neural network house: An environment that adapts to its inhabitants. In *Proc. AAAI Spring Symp. Intelligent Environments* (pp. 110–114). Retrieved from http://www.aaai.org/Papers/Symposia/Spring/1998/SS-98-02/SS98-02-017.pdf

Mozer, M. (1999). An intelligent environment must be adaptive. *IEEE Intelligent Systems and Their Applications*, *14*(2), 11–13. http://doi.org/10.1109/MIS.1999.757623

Mozer, M., Vidmar, L., & Dodier, R. (1997). The Neurothermostat: Predictive Optimal Control of Residential Heating Systems. In M. C. Mozer, M. I. Jordan, & T. Petsche (Eds.), *Advances in Neural Information Processing Systems 9* (pp. 953–959). Cambridge, MA: MIT Press. Retrieved from http://papers.nips.cc/paper/1299-the-neurothermostat-predictive-optimal-control-of-residential-heating-systems.pdf

Murphy, K. (2012). *Machine learning : A probabilistic perspective*. Cambridge, MA: MIT Press.

Noury, N., Berenguer, M., Teyssier, H., Bouzid, M.-J., & Giordani, M. (2011). Building an index of Activity of inhabitants from their activity on the Residential Electrical Power Line. *IEEE Transactions on Information Technology in Biomedicine*, *15*(5), 758–66. http://doi.org/10.1109/TITB.2011.2138149

Noury, N., Quach, K., Berenguer, M., Teyssier, H., Bouzid, M., Goldstein, L., & Giordani, M. (2011). Use of Electrical Devices Reveals Our Well Being. In *Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBC, 2011* (pp. 1769–1772). http://doi.org/10.1109/IEMBS.2011.6090505

Noury, N., Virone, G., Barralon, P., Ye, J., Rialle, V., & Demongeot, J. (2003). New trends in health smart homes. In *Proceedings 5th International Workshop on Enterprise Networking and Computing in Healthcare Industry 2003* (pp. 118–127). http://doi.org/10.1109/HEALTH.2003.1218728

Oppenheim, A. N. (2000). *Questionnaire design, interviewing and attitude measurement* (New ed.). London: Continuum.

P3 International. (2015). Kill A Watt. Retrieved January, 2015 from http://www.p3international.com/products/p4400.html

Paradiso, R. (2003). Wearable health care system for vital signs monitoring. In *4th International IEEE EMBS Special Topic Conference on Information Technology Applications in Biomedicine, 2003.* (pp. 283–286). IEEE. http://doi.org/10.1109/ITAB.2003.1222533

Park, S., & Jayaraman, S. (2003). Enhancing the quality of life through wearable technology. *IEEE Engineering in Medicine and Biology Magazine*, *22*(3), 41–48. http://doi.org/10.1109/MEMB.2003.1213625

Patel, S. N., Robertson, T., Kientz, J. A., Reynolds, M. S., & Abowd, G. D. (2007).

At the flick of a switch: Detecting and classifying unique electrical events on the residential power line. In J. Krumm, G. Abowd, A. Seneviratne, & T. Strang (Eds.), *UbiComp 2007: Ubiquitous Computing* (pp. 271–288). Springer Berlin Heidelberg. http://doi.org/10.1007/978-3-540-74853-3_16

Patten, M. L. (2007). *Understanding research methods : An overview of the essentials* (6th ed.). Glendale, CA: Pyrczak Publishing.

Percival, J., & Hanson, J. (2006). Big brother or brave new world? Telecare and its implications for older people's independence and social inclusion. *Critical Social Policy*, *26*(4), 888–909. http://doi.org/10.1177/0261018306068480

Perry, J., Beyer, S., & Holm, S. (2009). Assistive technology, telecare and people with intellectual disabilities: ethical considerations. *Journal of Medical Ethics*, *35*(2), 81–6. http://doi.org/10.1136/jme.2008.024588

Petersen, D., Steele, J., & Wilkerson, J. (2009). WattBot: a residential electricity monitoring and feedback system. In *Proceedings of the 27th international conference extended abstracts on Human factors in computing systems* (pp. 2847–2852). ACM. http://doi.org/10.1145/1520340.1520413

Pickard, A. (2013). *Research methods in information* (2nd ed.). London: Facet.

Piramuthu, S. (2004). Evaluating feature selection methods for learning in data mining applications. *European Journal of Operational Research*, *156*(2), 483–494. http://doi.org/10.1016/S0377-2217(02)00911-6

Pounds-Cornish, A., & Holmes, A. (2002). The iDorm-A Practical Deployment of Grid Technology. In *Proceedings of the 2nd IEEE/ACM International Symposium on Cluster Computing and the Grid, 2002* (p. 470). IEEE. http://doi.org/10.1109/CCGRID.2002.1017192

Punch, K. (2005). *Introduction to social research: quantitative and qualitative approaches* (2nd ed.). London: SAGE Publications.

Rahimi, S., Chan, A. D. C., & Goubran, R. A. (2011). Usage Monitoring of Electrical Devices in a Smart Home. In *2011 Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBC* (pp. 5307 – 5310). IEEE. http://doi.org/10.1109/IEMBS.2011.6091313

Rahimi, S., Chan, A. D. C., & Goubran, R. A. (2012). Nonintrusive Load Monitoring of Electrical Devices in Health Smart Homes. In *2012 IEEE International Instrumentation and Measurement Technology Conference (I2MTC)* (pp. 2313–2316). IEEE. http://doi.org/10.1109/I2MTC.2012.6229453

Reeder, B., Meyer, E., Lazar, A., Chaudhuri, S., Thompson, H. J., & Demiris, G. (2013). Framing the evidence for health smart homes and home-based consumer health technologies as a public health intervention for independent aging: a systematic review. *International Journal of Medical Informatics*, *82*(7), 565–79. http://doi.org/10.1016/j.ijmedinf.2013.03.007

Rogers, S., & Girolami, M. (2012). *A first course in machine learning*. Boca Raton: CRC Press.

Ruzzelli, A. G., Nicolas, C., Schoofs, A., & O'Hare, G. M. P. (2010). Real-Time Recognition and Profiling of Appliances through a Single Electricity Sensor. In *7th Annual IEEE Communications Society Conference on Sensor Mesh and Ad Hoc Communications and Networks (SECON), 2010* (pp. 1–9). IEEE. http://doi.org/10.1109/SECON.2010.5508244

Sathya, R., & Abraham, A. (2013). Comparison of Supervised and Unsupervised Learning Algorithms for Pattern Classification. *International Journal of Advanced Research in Artificial Intelligence*, *2*(2), 34–38. http://doi.org/10.14569/IJARAI.2013.020206

Sharma, B. (2010). Postpositivism. In A. J. Mills, G. Durepos, & E. Wiebe (Eds.), *Encyclopedia of Case Study Research.* (pp. 702–704). Thousand Oaks, CA: SAGE Publications. http://doi.org/10.4135/9781412957397

Sheik, S. S., Aggarwal, S. K., Poddar, A., Balakrishnan, N., & Sekar, K. (2004). A fast pattern matching algorithm. *Journal of Chemical Information and Computer Sciences*, *44*(4), 1251–1256. http://doi.org/10.1021/ci030463z

Singla, G., Cook, D. J., & Schmitter-Edgecombe, M. (2008). Incorporating Temporal Reasoning into Activity Recognition for Smart Home Residents. In *Proceedings of the AAAI Workshop on Spatial and Temporal Reasoning, 2008* (pp. 53–61). Retrieved from http://www.aaai.org/Papers/Workshops/2008/WS-08-11/WS08-11-008.pdf

Sintoni, A., Schoofs, A., Doherty, A., Smeaton, A. F., O'Hare, G., & Ruzzelli, A. G. (2011). Generating power footprints without appliance interaction: an enabler for privacy intrusion. In *2011 International Conference on Distributed Computing in Sensor Systems and Workshops (DCOSS),* (pp. 1–6). IEEE. http://doi.org/10.1109/DCOSS.2011.5982180

Sixsmith, A. (2000). An evaluation of an intelligent home monitoring system. *Journal of Telemedicine and Telecare*, *6*(2), 63–72. http://doi.org/10.1258/1357633001935059

Sixsmith, A., Hine, N., Neild, I., Clarke, N., Brown, S., & Garner, P. (2007). Monitoring the Well-being of Older People. *Topics in Geriatric Rehabilitation*, *23*(1), 9–23. Retrieved from http://journals.lww.com/topicsingeriatricrehabilitation/Fulltext/2007/01000/Monitoring_the_Well_being_of_Older_People.4.aspx

Smart Home Supplies. (n.d.).Retrieved January 2015 from http://www.smarthomesupplies.com/

Smarthome. (2015).Retrieved January 2015 from https://www.smarthome.com/

Song, Y., Morency, L. P., & Davis, R. (2013). Distribution-sensitive learning for imbalanced datasets. In *10th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition, FG, 2013* (pp. 1–6). IEEE. http://doi.org/10.1109/FG.2013.6553715

Soria, D., Garibaldi, J. M., Ambrogi, F., Biganzoli, E. M., & Ellis, I. O. (2011). A "non-parametric" version of the naive Bayes classifier. *Knowledge-Based Systems*, *24*(6), 775–784. http://doi.org/10.1016/j.knosys.2011.02.014

Spiegel, S., & Albayrak, S. (2014). Energy Disaggregation meets Heating Control. In *Proceedings of the 29th Annual ACM Symposium on Applied Computing (SAC '14)* (pp. 559–566). New York: ACM. http://doi.org/10.1145/2554850.2555088

Stone, J. (2013). *Bayes' rule : a tutorial introduction to Bayesian analysis*. Lexington, Kentucky: Sebtel Press.

Stowe, S., & Harding, S. (2010). Telecare, telehealth and telemedicine. *European Geriatric Medicine*, *1*(3), 193–197. http://doi.org/10.1016/j.eurger.2010.04.002

Tamura, T., Kawarada, A., Nambu, M., Tsukada, A., Sasaki, K., & Yamakoshi, K.-I. (2007). E-healthcare at an experimental welfare techno house in Japan. *The Open Medical Informatics Journal*, *1*, 1–7. http://doi.org/10.2174/1874431100701010001

Tamura, T., Togawa, T., Ogawa, M., & Yoda, M. (1998). Fully automated health monitoring system in the home. *Medical Engineering & Physics*, *20*(8), 573–9. http://doi.org/10.1016/S1350-4533(98)00064-2

Tan, P., Steinbach, M., & Kumar, V. (2006). *Introduction to data mining* (Internatio). Boston, Mass.; London: Pearson Addison Wesley.

Tang, P., & Venables, T. (2000). “Smart” homes and telecare for independent living. *Journal of Telemedicine and Telecare*, *6*(1), 8–14. http://doi.org/10.1258/1357633001933871

Tang, Y., Zhang, Y., Chawla, N. V., & Krasser, S. (2009). SVMs Modeling for Highly Imbalanced Classification. *IEEE Transactions on Systems, Man, and Cybernetics, Part B*, *39*(1), 281–288. http://doi.org/10.1109/TSMCB.2008.2002909

Tapia, E., Intille, S., & Larson, K. (2004). Activity Recognition in the Home Using Simple and Ubiquitous Sensors. In A. Ferscha & F. Mattern (Eds.), *Pervasive Computing* (Vol. 3001, pp. 158–175). Springer Berlin Heidelberg. http://doi.org/10.1007/978-3-540-24646-6_10

The Owl. (2015). OWL+USB. Retrieved January 2015 from http://www.theowl.com/index.php/energy-monitors/standalone-monitors/owl-usb/

Theodoridis, S., & Koutroumbas, K. (2009). *Pattern recognition* (4th ed.). Amsterdam ; London: Academic Press.

Theodoridis, S., Pikrakis, A., Koutroumbas, K., & Cavouras, D. (2010). *Introduction to Pattern Recognition: A MATLAB Approach*. Burlington, MA: Academic Press.

United Nations Department Of Economic And Social Affairs Population Division. (2013). *World Population Prospects: The 2012 Revision, Highlights and Advance Tables (Working Paper No. ESA/P/WP.228)*. Retrieved from http://esa.un.org/unpd/wpp/Documentation/pdf/WPP2012_HIGHLIGHTS.pdf

Virone, G., Noury, N., & Demongeot, J. (2002). A system for automatic measurement of circadian activity deviations in telemedicine. *IEEE Transactions on Biomedical Engineering*, *49*(12), 1463–1469. http://doi.org/10.1109/TBME.2002.805452

Walliman, N. (2006). *Social Research Methods*. London: SAGE Publications. Retrieved from https://www.dawsonera.com/abstract/9781847878182

Walliman, N. (2011). *Research methods: The basics*. London: Routledge. Retrieved from https://www.dawsonera.com/abstract/9780203836071

Wang, Z., Seidel, H. P., & Weinkauf, T. (2016). Multi-field Pattern Matching based on Sparse Feature Sampling. *IEEE Transactions on Visualization and Computer Graphics*, *22*(1), 807–816. http://doi.org/10.1109/TVCG.2015.2467292

Willmott, H. (2008). Critical theory. In R. Thorpe & R. Holt (Eds.), *The SAGE Dictionary of Qualitative Management Research* (pp. 67–69). London, United Kingdom: SAGE Publications.

Winett, R. A., Neale, M. S., & Grier, H. C. (1979). Effects of self-monitoring and feedback on residential electricity consumption. *Journal of Applied Behavior Analysis*, *12*(2), 173–184. http://doi.org/10.1901/jaba.1979.12-173

Witten, I., Frank, E., & Hall, M. (2011). *Data mining : practical machine learning tools and techniques.* (3rd ed.). Amsterdam: London: Morgan Kaufmann.

Yamazaki, T. (2005). Ubiquitous home: real-life testbed for home context-aware service. In *First International Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities, 2005* (pp. 54–59). http://doi.org/10.1109/TRIDNT.2005.37

Yamazaki, T. (2006). Beyond the smart home. In *International Conference on Hybrid Information Technology, 2006. ICHIT'06* (Vol. 2, pp. 350–355). IEEE. http://doi.org/10.1109/ICHIT.2006.253633

Yamazaki, T. (2007). The ubiquitous home. *International Journal of Smart Home*, *1*(1), 17–22. Retrieved from http://www.sersc.org/journals/IJSH/vol1_no1_2007/IJSH-2007-01-01-03.pdf

Yang, J., Wang, J., & Chen, Y. (2008). Using acceleration measurements for activity recognition : An effective learning algorithm for constructing neural classifiers. *Pattern Recognition Letters*, *29*(16), 2213–2220. http://doi.org/10.1016/j.patrec.2008.08.002

Yates, S. (2004). *Doing social science research*. London; Thousand Oaks, CA: Sage Publications in association with The Open University.

Youngblood, M., Cook, D. J., & Holder, L. B. (2005b). Seamlessly Engineering a Smart Environment. In *2005 IEEE International Conference on Systems, Man and Cybernetics* (Vol. 1, pp. 548–553). IEEE. http://doi.org/10.1109/ICSMC.2005.1571203

Youngblood, M., Cook, D., & Holder, L. (2005a). Managing Adaptive Versatile environments. *Pervasive and Mobile Computing*, *1*(4), 373–403. http://doi.org/10.1016/j.pmcj.2005.08.004

Zeifman, M., & Roth, K. (2011). Nonintrusive appliance load monitoring: Review and outlook. *IEEE Transactions on Consumer Electronics*, *57*(1), 76–84. http://doi.org/10.1109/TCE.2011.5735484

Zoha, A., Gluhak, A., Imran, M. A., & Rajasegarar, S. (2012). Non-intrusive load monitoring approaches for disaggregated energy sensing: a survey. *Sensors*, *12*(12), 16838–66. http://doi.org/10.3390/s121216838

Appendix One- Ethical Approval (Survey)

The University Of Sheffield.

Department Of Computer Science.

Head of Department
Professor John Derrick

Regent Court
211 Portobello Street
Sheffield
S1 4DP

Telephone: *+44 (0) 114 2221815*
Fax: *+44 (0) 114 2221810*
Email: *alice.tucker@dcs.shef.ac.uk*

15 October 2013

Jennifer Salter
Computer Science

Dear Jennifer

PROJECT TITLE: A QUESTIONNAIRE INTO NON-INTRUSIVE REMOTE ACTIVITY MONITORING OF AN ELDERLY/ILL RELATIVE

On behalf of the University ethics reviewers who reviewed your project, I am pleased to inform you that on 18 September 2012 the above-named project was unconditionally **approved** on ethics grounds, on the basis that you will adhere to the following document that you submitted for ethics review:

- University research ethics application form (*dated 22nd August 2012)*
- Participant information sheet (*dated 22nd Augusty 2012*)

If during the course of the project you need to deviate significantly from the above-approved document please inform me since written approval will be required. Please also inform me should you decide to terminate the project prematurely.

Yours sincerely

AB Tucker

Alice Tucker
Ethics Administrator

Appendix Two- Survey

A Questionnaire into Non-Intrusive Remote Activity Monitoring of an elderly or

We are conducting research into non-intrusive ways of monitoring daily activities in elderly or ill people and would like to assess the way, and to what extent a relative would wish to use this information.

What is your gender? *

- ◯ Male
- ◯ Female

What is your age? *

Do you currently care for an elderly or ill relative? *

- ◯ Yes
- ◯ No

Page 2 After page 1 Continue to next page

Note: "Go to page" selections will override this navigation. Learn more.

Have you ever cared for an elderly or ill relative? *

- ◯ Yes
- ◯ No

Page 3 After page 2 Continue to next page

Note: "Go to page" selections will override this navigation. Learn more.

What was the gender of your elderly or ill relative? *

- ◯ Male
- ◯ Female

Did your relative have a long term illness or disabling condition? *

- Yes
- No

Did your relative live alone? *

- Yes
- No

How far away from you did your relative live? *

- Within 1 mile
- 2-10 miles away
- 11-100 miles away
- More than 100 miles but within the UK
- In a different country

Page 4 | After page 3 | Go to page 5

If you have more than one elderly or ill relative, please answer these questions about the one which you are concerned about most.

What is the gender of your elderly or ill relative? *

- Male
- Female

What is your relative's age? *

Does your relative have a long term illness or disabling condition? *

- Yes
- No

Does your relative live alone? *

- Yes
- No

How far away from you does your relative live? *

- Within 1 mile
- 2-10 miles away
- 11-100 miles away
- More than 100 miles but within the UK
- In a different country

Page 5 After page 4 Continue to next page

The questions below are written for people who are currently caring for an elderly or ill relative. If that is not the case, please answer the questions as if you were caring for an elderly or ill relative.

Please list up to 3 events that most concern you, that may happen to your relative when they are alone:

Please list up to 3 activities that would, if you knew that they had been undertaken, give you assurance of your relative's current status:

Please rate how important you feel it would, be to know that your relative has undertaken these activities: *

	Very Important	Quite Important	Not at all Important
Changes in night time behaviour			
Waking up			
Food preparation			
Movement around the house			
Daytime general activities			

Please write down any other activities that you think it would be important to know about?

Please choose from the list the most important activity to be told that your relative has done: *

Please choose from the list the least important activity to be told that your relative has done: *

Page 6 After page 5 Continue to next page

Would you find it helpful to know that your relative has performed specific activities? For example:

That they have used the kettle? *

- Yes
- No

That they have watched TV *

- Yes
- No

That they have used the oven? *

- Yes
- No

That they have used a washing machine? *

- Yes
- No

That they have taken their medication? *

- Yes

- No

Please write down any other specific activities undertaken by your relative that will useful to know?

If there was a remote monitoring system in your relative's house, which type of activities would you find most relevant to be told about your relative: *

- Specific activities (e.g.that they turned the kettle on)
- General activities (e.g.that they are moving around the house)
- Both general and specific activities

In your opinion, what properties does a remote monitoring system need to have?

In your opinon, how important is it for a remote monitoring system to be non-intrusive? *

- Very important
- Quite important
- Not at all important

If there were an ability to provide non intrusive monitoring, with full agreement from your relative, who do you think should have access to this monitoring data? *

- Carers and close relatives who have a caring role
- All relavent health care profesionals
- Social services
- Other:

Thank you for you time!

Appendix Three- Ethics Approval and Data Sheet (Electricity data collection)

Information School Research Ethics Panel

Letter of Approval

Date: 11th July 2013

TO: Jennifer Salter

The Information School Research Ethics Panel has examined the following application:

Title: Understanding Activities Through the Analysis of Electricity Consumption Data

Submitted by: Jennifer Salter

And found the proposed research involving human participants to be in accordance with the University of Sheffield's policies and procedures, which include the University's '*Financial Regulations*', '*Good Research Practice Standards'* and the '*Ethics Policy Governing Research Involving Human Participants, Personal Data and Human Tissue'* (Ethics Policy).

This letter is the official record of ethics approval by the School, and should accompany any formal requests for evidence of research ethics approval.

Effective Date: 11th July 2013

Dr Angela Lin

Research Ethics Coordinator

The University of Sheffield. Information School	Understanding Activities Through the Analysis of Electricity Consumption Data

Researchers

Jennifer Salter (jasalter@sheffield.ac.uk) and Professor Peter Bath (p.a.bath@sheffield.ac.uk)

Purpose of the research

The aim of the research is to collect data on your electricity consumption from a series of households. The data from each household will be used to run a series of analyses, using a recognition model on the collected data, to test whether (using changes to electrical consumption) the person/people within the house have performed certain specific and/or general activities or tasks. The proposed length for you to take part in this study is 1 week.

Who will be participating?

Close family and friends of the researchers. Consent for this study will be obtained from every member of the household over the age of 18. If your household contains people under the age of 18, you will need to give consent on their behalf.

What will you be asked to do?

In participating in the research there is no need to change your daily activities but to allow electricity consumption data to be collected and reviewed, you are also asked to complete a date and time diary of when certain appliances are used and/or you have done certain activities (the diary will be provided by me). An example of the list of activities and appliances can be found in appendix 2.

What are the potential risks of participating?

The risks of participating are the same as those experienced in everyday life, although you might find it slightly inconvenient to record these events.

What data will we collect?

The collection of the data will have no affect on your electricity consumption nor will it require any alteration to electrical cabling in the house. The data will be collected through an energy monitor placed as outlined below inside your house (a picture of the equipment is in appendix 1). The data is collected by:

- Placing a mains sensor- which clips around the main electrical feed cable into the house consumer unit (the box with all the fuses in)
- An energy monitor that provides readings on the electricity consumption used in the house.
- A data logger, which need to be plugged into a socket and will store the electricity consumption and download data to a secure web server when connected to the internet

Once the equipment is installed the equipment will automatically send/store the electricity consumption for the house. In addition, I am asking you to record when you have used particular appliances, and undertaken certain activities.

What will we do with the data?

This data is needed to build, test and validate a model used to recognise certain specific and/ or general activities or tasks from a house's electricity consumption. The data collected will be stored on a secure University of Sheffield server. The data collected will be only used as specified in the aims of the research and can be deleted from this server at any time requested by the participant. The data is received with an ID number that cannot identify the household that it is related to and only the researcher will know which ID's belong to which household. The data you provide will only be used for this project and will be destroyed at the end of it.

Will my participation be confidential?

All information collected, as part of this research will be keep strictly confidential. The electricity consumption data and the diary are both coded with unique ID numbers that will not identify the participants or the location of the house where the data comes from. The researcher will know, both from the diary and electrical consumption data the times certain appliances are switched on /off. Individual participant will not be identified in any reports or publications drawn from this research.

What will happen to the results of the research project?

The results collected from this study will be used for the completion of the main researcher's PhD. Results from the study may also be disseminated in research papers.

I confirm that I have read and understand the description of the research project, and that I have had an opportunity to ask questions about the project.

I understand that my participation is voluntary and that I am free to withdraw at any time without any negative consequences.

I understand that I may decline to answer any particular question or questions, or to do any of the activities. If I stop participating at all time, all of my data will be purged.

I understand that my responses will be kept strictly confidential, that my name or identity will not be linked to any research materials, and that I will not be identified or identifiable in any report or reports that result from the research.

I give permission for the research team members to have access to my anonymised responses.

I give permission for the research team to re-use my data for future research as specified above.

I agree to take part in the research project as described above.

Participant Name (Please print)	Participant Signature
Participant Name (Please print)	Participant Signature
Participant Name (Please print)	Participant Signature
Participant Name (Please print)	Participant Signature
Researcher Name (Please print)	Researcher Signature

Date

Note: If you have any difficulties with, or wish to voice concern about, any aspect of your participation in this study, please contact Dr. Angela Lin, Research Ethics Coordinator, Information School, The University of Sheffield (ischool_ethics@sheffield.ac.uk), or to the University Registrar and Secretary.

Appendix 1 – Picture of the equipment

The electricity monitor and the mains sensor

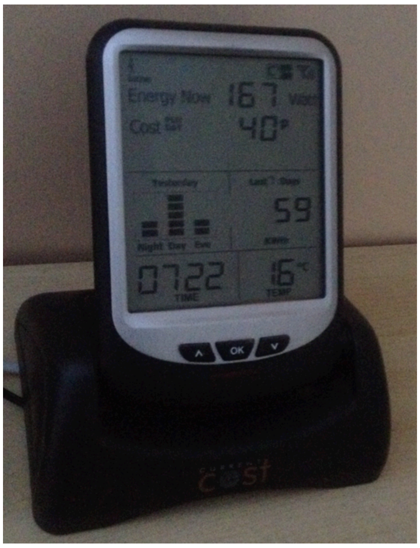

The data logger

Appendix 2 – Ideas of appliances to be monitored

Examples of what appliances that you will need to keep a diary of what you use is shown below. (NOTE: some of these appliances might not be applicable to you; it depends if you have each one in your home and/or if they are gas or electric.

Examples of appliances:

- Kettle
- Oven
- Hobs
- Television
- Washing machine
- Dishwasher
- Toaster
- Electric shower
- Microwave

Other appliances might be added to the list, depending on how many of theses electrical appliances above you have in your home.

Examples of activities:

- Food preparation
- Waking up

Appendix Four- Results From Trial Analysis Trials 1-4

Window size	Appliance	True Positive	False Positive	False Negative	True Negative	Sensitivity	Specificity	PPV	NPV	Overall PPV	Overall Sensitivity
3	Shower	3	2	0	39281	100%	99.995%	60%	100%	0.75%	72.73%
	Microwave	3	14	1	39268	75%	99.964%	17.647%	99.997%		
	Kettle	1	1000	1	38284	50%	97.454%	0.100%	99.997%		
	Dishwasher	1	41	1	39243	50%	99.896%	2.381%	99.997%		
4	Shower	3	2	0	39281	100%	99.995%	60%	100%	0.71%	72.73%
	Microwave	3	7	1	39275	75%	99.982%	30%	99.997%		
	Kettle	1	1068	1	38216	50%	97.281%	0.094%	99.997%		
	Dishwasher	1	42	1	39242	50%	99.893%	2.326%	99.997%		
5	Shower	3	2	0	39281	100%	99.995%	60%	100%	1.42%	72.73%
	Microwave	3	3	1	39279	75%	99.992%	50%	99.997%		
	Kettle	1	504	1	38780	50%	98.717%	0.198%	99.997%		
	Dishwasher	1	45	1	39239	50%	99.885%	2.174%	99.997%		
6	Shower	3	2	0	39281	100%	99.995%	60%	100%	**1.80%**	**72.73%**
	Microwave	3	3	1	39279	75%	99.992%	50%	99.997%		
	Kettle	1	384	1	38900	50%	99.023%	0.260%	99.997%		
	Dishwasher	1	48	1	39236	50%	99.878%	2.041%	99.997%		
7	Shower	3	2	0	39281	100%	99.995%	60%	100%	1.69%	72.73%
	Microwave	3	3	1	39279	75%	99.992%	50%	99.997%		
	Kettle	1	412	1	38872	50%	98.951%	0.242%	99.997%		
	Dishwasher	1	48	1	39236	50%	99.878%	2.041%	99.997%		
8	Shower	3	2	0	39281	100%	99.995%	60%	100%	1.48%	72.73%
	Microwave	3	3	1	39279	75%	99.992%	50%	99.997%		
	Kettle	1	481	1	38803	50%	98.776%	0.207%	99.997%		
	Dishwasher	1	48	1	39236	50%	99.878%	2.041%	99.997%		
9	Shower	3	2	0	39281	100%	99.995%	60%	100%	1.05%	63.64%
	Microwave	3	4	1	39278	75%	99.990%	42.857%	99.997%		
	Kettle	1	522	1	38762	50%	98.671%	0.191%	99.997%		
	Dishwasher	0	132	2	39152	0%	99.664%	0%	99.995%		
10	Shower	3	2	0	39281	100%	99.995%	60%	100%	0.75%	63.64%
	Microwave	3	11	1	39271	75%	99.972%	21.429%	99.997%		
	Kettle	1	572	1	38712	50%	98.544%	0.175%	99.997%		
	Dishwasher	0	342	2	38942	0%	99.129%	0%	99.995%		

Table A4.1: The results from the first attempt for window size 3 to 10

Window size	Appliance	True Positive	False Positive	False Negative	True Negative	Sensitivity	Specificity	PPV	NPV	Overall PPV	Overall Sensitivity
3	Shower	3	0	0	39283	100%	100%	100%	100%	1.31%	72.73%
	Microwave	3	108	1	39174	75%	99.725%	2.703%	99.997%		
	Kettle	2	492	0	38792	100%	98.748%	0.405%	100%		
	Dishwasher	0	2	2	39282	0%	99.995%	0%	99.995%		
4	Shower	3	0	0	39283	100%	100%	100%	100%	1.59%	72.73%
	Microwave	3	11	1	39271	75%	99.972%	21.429%	99.997%		
	Kettle	1	454	1	38830	50%	98.844%	0.220%	99.997%		
	Dishwasher	1	29	1	39255	50%	99.926%	3.333%	99.997%		
5	Shower	3	0	0	39283	100%	100%	100%	100%	1.39%	72.73%
	Microwave	3	9	1	39273	75%	99.977%	25%	99.997%		
	Kettle	1	526	1	38758	50%	98.661%	0.190%	99.997%		
	Dishwasher	1	32	1	39252	50%	99.919%	3.030%	99.997%		
6	Shower	3	0	0	39283	100%	100%	100%	100%	3.16%	72.73%
	Microwave	3	9	1	39273	75%	99.977%	25%	99.997%		
	Kettle	1	204	1	39080	50%	99.481%	0.488%	99.997%		
	Dishwasher	1	32	1	39252	50%	99.919%	3.030%	99.997%		
7	Shower	3	0	0	39283	100%	100%	100%	100%	5.56%	72.73%
	Microwave	3	6	1	39276	75%	99.985%	33.333%	99.997%		
	Kettle	1	90	1	39194	50%	99.771%	1.099%	99.997%		
	Dishwasher	1	40	1	39244	50%	99.898%	2.439%	99.997%		
8	Shower	3	0	0	39283	100%	100%	100%	100%	**5.97%**	**72.73%**
	Microwave	3	4	1	39278	75%	99.990%	42.857%	99.997%		
	Kettle	1	81	1	39203	50%	99.794%	1.220%	99.997%		
	Dishwasher	1	41	1	39243	50%	99.896%	2.381%	99.997%		
9	Shower	3	0	0	39283	100%	100%	100%	100%	3.95%	90.91%
	Microwave	4	29	0	39253	100%	99.926%	12.121%	100%		
	Kettle	1	29	1	39255	50%	99.926%	3.333%	99.997%		
	Dishwasher	2	185	0	39099	100%	99.529%	1.070%	100%		
10	Shower	3	0	0	39283	100%	100%	100%	100%	2.24%	72.73%
	Microwave	3	64	1	39218	75%	99.837%	4.478%	99.997%		
	Kettle	0	20	2	39264	0%	99.949%	0%	99.995%		
	Dishwasher	2	265	0	39019	100%	99.325%	0.749%	100%		

Table A4.2: The results from trial analysis 3 without filter for window size 3 to

Window size	Appliance	True Positive	False Positive	False negative	True negative	Sensitivity	Specificity	PPV	NPV	Overall PPV	Overall Sensitivity
3	Shower	3	3	0	31384	100%	99.990%	50%	100%	4.08%	72.73%
	Microwave	3	112	1	31274	75%	99.643%	2.609%	99.997%		
	Kettle	2	68	0	31320	100%	99.783%	2.857%	100%		
	Dishwasher	0	5	2	31383	0%	99.984%	0%	99.994%		
4	Shower	3	6	0	31791	100%	99.981%	33.333%	100%	4.62%	72.73%
	Microwave	3	44	1	31752	75%	99.862%	6.383%	99.997%		
	Kettle	2	48	0	31750	100%	99.849%	4%	100%		
	Dishwasher	0	67	2	31731	0%	99.789%	0%	99.994%		
5	Shower	3	9	0	32337	100%	99.972%	25%	100%	3.42%	72.73%
	Microwave	3	43	1	32302	75%	99.867%	6.522%	99.997%		
	Kettle	2	134	0	32213	100%	99.586%	1.471%	100%		
	Dishwasher	0	40	2	32307	0%	99.876%	0%	99.994%		
6	Shower	3	12	0	32671	100%	99.963%	20%	100%	**3.66%**	**81.82%**
	Microwave	3	44	1	32638	75%	99.865%	6.383%	99.997%		
	Kettle	2	88	0	32596	100%	99.731%	2.222%	100%		
	Dishwasher	1	93	1	32591	50%	99.715%	1.064%	99.997%		
7	Shower	3	14	0	33060	100%	99.958%	17.647%	100%	3.11%	72.73%
	Microwave	3	51	1	33022	75%	99.846%	5.556%	99.997%		
	Kettle	2	98	0	32977	100%	99.704%	2%	100%		
	Dishwasher	0	86	2	32989	0%	99.740%	0%	99.994%		
8	Shower	3	17	0	33248	100%	99.949%	15%	100%	2.80%	72.73%
	Microwave	3	52	1	33212	75%	99.844%	5.455%	99.997%		
	Kettle	2	165	0	33101	100%	99.504%	1.198%	100%		
	Dishwasher	0	44	2	33222	0%	99.868%	0%	99.994%		
9	Shower	3	19	0	33653	100%	99.944%	13.636%	100%	1.41%	54.55%
	Microwave	0	118	4	33553	0%	99.650%	0%	99.988%		
	Kettle	1	78	1	33595	50%	99.768%	1.266%	99.997%		
	Dishwasher	2	205	0	33468	100%	99.391%	0.966%	100%		
10	Shower	3	21	0	33919	100%	99.938%	12.5%	100%	1.10%	54.55%
	Microwave	0	214	4	33725	0%	99.369%	0%	99.988%		
	Kettle	1	82	1	33859	50%	99.758%	1.205%	99.997%		
	Dishwasher	2	221	0	33720	100%	99.349%	0.897%	100%		

Table A4.3: The results from trial analysis 4 without filter for window size 3 to 10

Appendix Five

Combo	Feature set		
1	avg	peak	std
2	avg	peak	rms
3	avg	peak	peak to avg
4	avg	peak	peak to rms
5	avg	peak	rms to avg
6	avg	std	rms
7	avg	std	peak to avg
8	avg	std	peak to rms
9	avg	std	rms to avg
10	avg	rms	peak to avg
11	avg	rms	peak to rms
12	avg	rms	rms to avg
13	avg	peak to avg	peak to rms
14	avg	peak to avg	rms to avg
15	avg	peak to rms	rms to avg
16	peak	std	rms
17	peak	std	peak to avg
18	peak	std	peak to rms
19	peak	std	rms to avg
20	peak	rms	peak to avg
21	peak	rms	peak to rms
22	peak	rms	rms to avg
23	peak	peak to avg	peak to rms
24	peak	peak to avg	rms to avg
25	peak	peak to rms	rms to avg
26	std	rms	peak to avg
27	std	rms	peak to rms
28	std	rms	rms to avg
29	std	peak to avg	peak to rms
30	std	peak to avg	rms to avg
31	std	peak to rms	rms to avg
32	rms	peak to avg	peak to rms
33	rms	peak to avg	rms to avg
34	rms	peak to rms	rms to avg
35	peak to avg	peak to rms	rms to avg

Combo	Feature set			
36	avg	peak	std	rms
37	avg	peak	std	peak to avg
38	avg	peak	std	peak to rms
39	avg	peak	std	rms to avg
40	avg	peak	rms	peak to avg
41	avg	peak	rms	peak to rms
42	avg	peak	rms	rms to avg
43	avg	peak	peak to avg	peak to rms
44	avg	peak	peak to avg	rms to avg
45	avg	peak	peak to rms	rms to avg
46	avg	std	rms	peak to avg
47	avg	std	rms	peak to rms
48	avg	std	rms	rms to avg
49	avg	std	peak to avg	peak to rms
50	avg	std	peak to avg	rms to avg
51	avg	std	peak to rms	rms to avg
52	avg	rms	peak to avg	peak to rms
53	avg	rms	peak to avg	rms to avg
54	avg	rms	peak to rms	rms to avg
55	avg	peak to avg	peak to rms	rms to avg
56	peak	std	rms	peak to avg
57	peak	std	rms	peak to rms
58	peak	std	rms	rms to avg
59	peak	std	peak to avg	peak to rms
60	peak	std	peak to avg	rms to avg
61	peak	std	peak to rms	rms to avg
62	peak	rms	peak to avg	peak to rms
63	peak	rms	peak to avg	rms to avg
64	peak	rms	peak to rms	rms to avg
65	peak	peak to avg	peak to rms	rms to avg
66	std	rms	peak to avg	peak to rms
67	std	rms	peak to avg	rms to avg
68	std	rms	peak to rms	rms to avg
69	std	peak to avg	peak to rms	rms to avg
70	rms	peak to avg	peak to rms	rms to avg

Combo	Feature set						
71	avg	peak	std	rms	peak to avg		
72	avg	peak	std	rms	peak to rms		
73	avg	peak	std	rms	rms to avg		
74	avg	peak	std	peak to avg	peak to rms		
75	avg	peak	std	peak to avg	rms to avg		
76	avg	peak	std	peak to rms	rms to avg		
77	avg	peak	rms	peak to avg	peak to rms		
78	avg	peak	rms	peak to avg	rms to avg		
79	avg	peak	rms	peak to rms	rms to avg		
80	avg	peak	peak to avg	peak to rms	rms to avg		
81	avg	std	rms	peak to avg	peak to rms		
82	avg	std	rms	peak to avg	rms to avg		
83	avg	std	rms	peak to rms	rms to avg		
84	avg	std	peak to avg	peak to rms	rms to avg		
85	avg	rms	peak to avg	peak to rms	rms to avg		
86	peak	std	rms	peak to avg	peak to rms		
87	peak	std	rms	peak to avg	rms to avg		
88	peak	std	rms	peak to rms	rms to avg		
89	peak	std	peak to avg	peak to rms	rms to avg		
90	peak	rms	peak to avg	peak to rms	rms to avg		
91	std	rms	peak to avg	peak to rms	rms to avg		
92	avg	peak	std	rms	peak to avg	peak to rms	
93	avg	peak	std	rms	peak to avg	rms to avg	
94	avg	peak	std	rms	peak to rms	rms to avg	
95	avg	peak	std	peak to avg	peak to rms	rms to avg	
96	avg	peak	rms	peak to avg	peak to rms	rms to avg	
97	avg	std	rms	peak to avg	peak to rms	rms to avg	
98	peak	std	rms	peak to avg	peak to rms	rms to avg	
99	avg	peak	std	rms	peak to avg	peak to rms	rms to avg

Table A5.1: Key of feature set combinations

Appendix Six- Appliance Attributes

	Peak	Standard deviation	Peak to average ratio
Microwave	1284	661.366	2.535
	1259	649.771	2.539
	1222	624.061	2.496
	1197	622.605	2.586
	1203	577.779	2.746
	1064	578.915	3.265
	1195	584.688	2.817
	826	377.674	5.326
	1225	632.064	2.574
	1264	649.356	2.555
Oven	2208	1088.883	2.666
	2112	1046.668	2.653
	2282	1228.200	2.933
	2047	1031.881	2.429
Shower	9169	4756.321	2.533
	8571	4347.201	2.481
	9095	4699.242	2.517
	9230	4722.747	2.528
	9036	4702.406	2.532
	9196	4726.872	2.546
	9204	4724.767	2.553
	9013	4660.288	2.512
Dishwasher	1822	940.030	2.510
	2007	1008.734	2.584
	3862	1832.598	2.886
	1993	972.791	2.531
Washing machine	1518	852.320	3.333
	609	353.568	4.250
	1697	817.711	3.040
	1205	681.605	3.148
	1648	730.835	3.503
	2082	817.944	5.635
	1608	799.234	3.010

Table A6.1: Attribute values for each class from trial house. For window size 6 backwards 4 forwards, feature set (peak, standard deviation and peak to average ratio)

	Average	Standard deviation	Root mean square	Peak to root mean square ratio
Kettle[1]	1138.681	1315.690	1611.817	1.439
Dishwasher	769.250	1266.119	1339.417	1.417
	1090.500	1257.603	1541.222	1.429
	1335.000	1202.454	1693.117	1.420
	1121.500	1289.710	1582.807	1.440
Oven	1252.000	1539.799	1829.136	1.426
	1298.250	1564.352	1876.393	1.419
	1314.500	1587.509	1902.117	1.434
	764.250	891.476	1086.335	1.441
Washing machine	1364.250	1086.663	1657.349	1.417
	1455.000	1123.722	1750.454	1.480
	1271.250	1288.593	1691.577	1.503
	1331.750	1268.757	1726.519	1.413
	1335.500	932.258	1560.573	1.434
	1533.500	933.090	1733.383	1.417
	1093.000	1213.963	1516.552	1.450

Table A6.2: Attribute values for each class from household one. For window size 2 backwards 2 forwards, feature set (average, standard deviation, root mean square and peak to root mean square ratio) [1]Due to the large number of kettle points an average is shown

	Average	Peak	Standard deviation	Root mean square	Peak to average ratio	Peak to root mean square ratio
Kettle[1]	2096.15	3010	1323.767	2437.731	1.439	1.232
Dishwasher	1817.50	2629	1224.118	2148.130	1.446	1.224
	1764.75	2469	1149.765	2066.654	1.399	1.195
	1838.38	2543	1138.924	2124.765	1.383	1.197
Toaster	522.50	718	320.976	602.622	1.374	1.191
	506.00	684	287.587	573.065	1.352	1.194
	373.50	887	362.805	504.655	2.375	1.758
Microwave	971.50	1388	636.497	1139.430	1.429	1.218
	458.38	966	654.356	764.700	2.107	1.263
	346.13	1002	687.612	730.419	2.895	1.372
	816.38	1326	585.593	983.119	1.624	1.349
	862.25	2551	1989.922	2051.412	2.959	1.244
	897.25	1445	702.737	1112.280	1.610	1.299

Table A6.3: Attribute values for each class from household 2. For window size 2 backwards 6 forwards, feature set (average, peak, standard deviation, root mean square, peak to average ratio, peak to root mean square ratio) [1]Due to the large number of kettle points an average is shown

	Peak	Standard deviation	Root mean square	Root mean square to average ratio
Kettle[1]	3052.545	1604.762	2223.734	0.636
Shower	8290	4350.838	6151.967	1.323
	7835	4130.067	5836.306	1.324
	8094	4166.888	5889.832	1.323
	8161	4288.063	6064.029	1.323
	7743	4060.973	5742.053	1.323
	7949	4186.487	5915.966	1.324
	7622	4163.735	5553.329	1.389
	8169	4314.992	6102.056	1.323
	6256	4135.077	4695.981	1.727
	8247	4352.695	6157.305	1.323
	8385	4239.826	5989.979	1.324
Washing machine	1740	975.024	1210.069	1.502
	1801	958.832	1291.251	1.377
	1833	990.526	1202.144	1.547
	2210	1045.113	1577.574	1.266
	2090	985.989	1484.654	1.268
	2252	1026.341	1656.021	1.221
	2460	1051.126	1752.581	1.202
	3573	1799.174	2284.981	1.461
Toaster (2)[1]	1261.533	597.792	884.852	1.313
Toaster (4)	1880	932.205	1412.143	1.263
	1927	1013.737	1433.521	1.323
	1951	1029.685	1457.657	1.322
	1937	1021.784	1444.121	1.323

Table A6.4: Attribute values for each class from household three. For window size 3 backwards 4 forwards, feature set (peak, standard deviation, root mean square, peak to average ratio, peak) [1]Due to the large number of points an average is shown