Master thesis
An intelligent monitoring system

Supporting mentally disabled mothers by monitoring object interactions

August 31, 2009

Marjolein Terhaard (s0105848)
m.terhaard@student.ru.nl
Radboud Universiteit Nijmegen

Supervisors:
dr. L.G. Vuurpijl (Radboud University Nijmegen)
dr. I.G. Sprinkhuizen-Kuyper (Radboud University Nijmegen)
B.D. Williams (CENSI)

Acknowledgements

There are a lot of people I want to thank because of the great deal of support they have provided during my internship and the accomplishment of this thesis. First of all I want to thank my supervisors Louis Vuurpijl and Ida Sprinkhuizen-Kuyper for giving me direction to my research and for their feedback on my evolving thesis.
I want to thank Hans Appel and Gea Koeling from CENSI for convincing me to do my internship at the CENSI, and Hans for his inspiring conversations especially and help during my internship.
Bryan Williams, also from CENSI, helped me a lot by thinking along with me when things got though during my internship and afterwards with providing feedback on my thesis. Georges Meinders, Twan Talens and Stefan Hemmes I want to thank for their help in connecting the sensors and making sure that my test environment worked and was safe to use. My fellow students at CENSI and the other people at the HIT for the nice time I have had.
My thanks also goes to Janneke van der Burgh and Rik Bakker from NOVO, for their information about my target group and inspiration for my research.
My dear friends Esther Ticheler, Desiree Spronk, Saskia Robben, Tom Schut, Ron van Kesteren and Jop van Heesch provided me with a lot of moral support. And my parents Theo and Margo Terhaard and Nils van der Lubbe I want to thank for always being there for me.

Introduction of CENSI

I have done my internship at CENSI. CENSI stands for 'Center of Excellence for Intelligent Sensor Innovation' and is a cooperation between the 'Hanze Hoge School Groningen', 'Hanze Institute of Technology' and 'Sun Microsystems'. It is specialized in the appliance of intelligent sensor systems. The CENSI guides students from different disciplines (for example informatics, human-technology, nursing) in projects commissioned by different companies such as TNO ICT, NOVO and NEDAP. All projects are dealing with sensors innovation. The CENSI is situated in the building of the 'Hanze Institute of Technology', which is specialized in innovative sensor technology and at the Hanze Hoge School Groningen.

Abstract

In this thesis the research on the design of an intelligent monitoring system to support people with a mental disability with their daily life routines is described. The focus has been on the automatic classification of preparing a bottle for a child. Relevant data was needed for classification, which was not yet available. We therefore built a test environment, consisting of different electronic devices and other objects used when preparing a bottle. These objects were equipped with sensors to detect when the objects were used. We asked subjects to perform several tasks in this test environment. A software platform was built to handle these data from the experiments. Different classification algorithms were tested on these data. The results showed that in 90% of the cases the tasks were correctly classified by the Support Vector Machine algorithm.

Contents

Chapter 1

Introduction

The research described in this thesis pursues the development of an intelligent monitoring system to support people with a mental disability with their daily life routines. An intelligent monitoring system is a system that autonomously monitors its environment through its sensors, processes the data from the sensors and reports about these data and/or takes actions. Processing of the data means for example that the data is automatically classified to diagnose problems, or that the system uses its reasoning capacities to plan actions. For the processing capabilities of intelligent monitoring systems Artificial Intelligence (AI) techniques are used, such as machine learning and logic. More about these techniques is desribed in Section 2.5.

There is a need for surveillance and inspection in many domains. In some of these domains the use of an intelligent monitoring system is required to assist or take over the task of surveillance and inspection from the human worker. Reasons for this can be that the human resources are scarce, too expensive or that the environment is too dangerous.

A domain where monitoring systems are used is that of environmental monitoring. For example to detect the exact origin of a forest fire by the use of of a network of sensor nodes. These sensors nodes are spread around the forest and are programmed to collaborate and automatically form a network that can detect if and where a fire starts (Akyildiz, Su, Sankarasubramaniam, & Cayirci, 2002). Another domain is that of healthcare, where the physical condition of patients is automatically monitored (Milenkovic, Otto, & Jovanov, 2006; Gao, Greenspan, Welsh, Juang, & Alm, 2005). Another example in the healthcare domain is the monitoring of the activities of people. This in order to determine if (elderly and/or disabled) people are still capable of taking care of themselves. In Section 2.3 we will go further into the different domains where intelligent monitoring systems are used.

Humans have fives senses that they use for monitoring, their eyes, ears, tongue, skin and nose. An intelligent monitoring system uses artificial sensors for monitoring. There are a lot of different sensors, mostly classified to the kind of quantity measured; mechanical quantities (force, flow rate and pressure), thermal quantities (temperature), electrostatic, magnetic fields, fluxes, radiation intensity (electromagnetic and nuclear intensities), chemical quantities (concentration of humidity and gas components), biological quantities (concentrations of enzyme substrates and antibodies) (Harsányi, Bojta, Gordon, Lepsényi, & Ballun, 2009). In Section 2.2 we will explain more about the sensors used in monitoring systems.

In this thesis we will describe an intelligent monitoring system that will be used to monitor

the activities of mentally disabled people to support them in their daily life.
In this introduction we will first describe why these mentally disabled people need this support and what the problems are in providing this support. We then will explain why an intelligent monitoring system could be a solution and we will introduce the research questions of this thesis.

1.1 Problem

Being safe and independent are values every human being longs for, including people with a mental disability. Because of the disorder of these people, these values are sometimes at odds. The precautions taken to make sure that people with a mental disability are safe, sometimes have a negative impact on their independence.
The DSM-IV[1] (American Psychiatric Association, 1994) uses the term Mental Retardation. This term is the medical term, and is commonly replaced by the term mental disability. The DSM-IV defines Mental Retardation as:

> The essential feature of Mental Retardation is significantly subaverage general intellectual functioning (Criterion A) that is accompanied by significant limitations in adaptive functioning in at least two of the following skill areas: communication, self-care, home living, social/ interpersonal skills, use of community resources, self-direction, functional academic skills, work, leisure, health, and safety (Criterion B). The onset must occur before the age 18 years (Criterion C).

People with a mild Mental Retardation usually achieve sufficient social and vocational capabilities for minimal self-support. Supervision, guidance and assistance still may be needed, especially when under unusual social or economic stress (American Psychiatric Association, 1994).
Problems with these people occur because, for example, they have trouble getting in and out of bed on time. They lack the overview and planning skills to see that they have to go to bed on time in order to be able to go to work on time the next day. Their day-night rhythm gets disturbed which can cause them to lose their job or to forget to take their medicines on time. To prevent these negative effects, support for people with a mental disability is provided by healthcare organisations. Employees of these healthcare organisations monitor the mentally disabled clients. They collect information about the client, such as what actions the client is performing, or what the healthcare status of the client is. They interpret this information and signal problems or improvements in, for example, the clients capability of taking care of her/himself.
The healthcare workers report their findings to the responsible persons and/or to the clients. When it is necessary, the healthcare workers will intervene, and take actions to prevent or solve problems (Bakker, 2008).

There are some drawbacks to the use of human resources:

- Human intervention and monitoring is very expensive and there is a shortage of staff that wants to work in the healthcare business

[1] Diagnostic and Statistical Manual of Mental Disorders

- It has a huge impact on the privacy and the autonomy of these clients
- The clients do not always accept the help, because they do not always agree that they have a problem
- Most of the time actions are only taken when things have gone wrong

A solution is therefore sought for these problems.

1.2 Solution

A solution could be to use an intelligent monitoring system which could then (partly) take over the tasks of monitoring, signaling, reporting and intervention that until now is done by healthcare workers. An intelligent monitoring system could solve the problems described above (Bakker, 2008):

- The system would take over some tasks of the human healthcare workers, the problem of a shortage of staff could be partially solved
- The system would watch continuously over the target person, so actions can be taken before something goes wrong
- The target person is monitored without the need of having healthcare workers in their home 24 hours a day. This will increase the experience of personal freedom of the person
- Constant feedback can be provided on the target person's actions, he/she could learn from this feedback, and this will contribute to his/her personal growth
- Help from a computer is probably more accepted especially by young clients, because they think computer systems are more interesting then human healthcare workers

An intelligent monitoring system should take over the task of monitoring, signaling, reporting and intervention. Monitoring can be done by the system by automatically gathering data from the environment through its sensors. The signaling is a task of the processing part of the system. Furthermore the reporting and intervention are tasks of the output of the system.
In this thesis we will focus on the first two tasks, monitoring and signaling. The question is: "Can an intelligent monitoring system be used for the monitoring of mentally disabled people and the signaling of problems".
We will answer this question by focussing on a small problem domain. The problems of mentally disabled people are very diverse and personal. To make an intelligent monitoring system to cover all these problems would exceed the scope of this thesis. Therefore we have chosen to center on one problem group, the group of mentally disabled mothers. Our focus will be on the problem that sometimes the mother is too preoccupied by herself and she forgets to prepare a bottle of powder milk for her baby (J. van der Burgh, teamcoordinator NOVO ambulant team, personal communication, October 8, 2008). The reason that we have chosen for this problem is explained in more detail in Chapter 3.
We want to propose an intelligent monitoring system that can recognize that a bottle of milk has been prepared for a child, and when this action has not been conducted for a certain time, a signal could be given to the mother to remember her that she has to feed her child.

We want to investigate how we can detect that a bottle has been prepared, so what sensors to use. This relates to the monitoring task of the system, namely the input of the system. Additionally we want to investigate how reliable it can be classified that a bottle has been prepared. This relates to the second task, signaling, which will be done by the processing part of the system. These research questions will be pursued in this thesis. They can be summarized as follows:

RQ1. How to monitor that a bottle has been prepared

RQ2. How reliable can it be classified that a bottle for the child has been prepared

In order to classify the task of preparing a bottle we needed a relevant dataset. Such a dataset was not available, so we decided to generate a dataset ourself. We built a test kitchen consisting of objects that were needed for our problem task of preparing a bottle. We equipped this kitchen with sensors to monitor the use of objects. Additionally we built a software platform to handle the data from the sensors. We have used this kitchen to gather relevant data from experiments with test subjects. The test subjects were asked to perform multiple tasks in the kitchen, including preparing a bottle. With the dataset obtained by these experiments we could test different classification algorithms.

The remainder of this thesis is structured as follows. In Chapter 2 we will first give a review of intelligent monitoring systems. We will elaborate on the requirements of these system and we will present some example domains. Furthermore it will be explained how behaviour of people can be monitored on the basis of object interaction. Followed by an overview of the different sensors and Artificial Intelligence techniques used in intelligent monitoring systems. Chapter 3 contains an overview of the research environment and experimental design. In this chapter the first research question (RQ1) of how to monitor that a bottle has been prepared, will be answered. In Chapter 4 we will present how well the machine learning algoritms were able to classify the tasks from the data of the experiments, thereby answering RQ2.
An assessment will be made in Chapter 5 about the extend to which RQ1 and RQ2 were answered in this thesis. Furthermore some general conclusions and future research will be described.

Chapter 2

Review on intelligent monitoring systems

In this chapter we will give a review on intelligent monitoring systems. Although extensive research into intelligent monitoring systems has been done, a clear definition of an intelligent monitoring system is never given. Therefore we will give our own definition of an intelligent monitoring system:

> An intelligent monitoring system is a system that that autonomously monitors its environment through its sensors, processes the data from the sensors and reports about these data and/or takes actions.

In Section 2.1 we will explain what the requirements are for an intelligent monitoring system, followed by an overview of the sensors used in intelligent monitoring systems (Section 2.2). In Section 2.3 some example domains are described that use intelligent monitoring systems. We will dive deeper into the domain of monitoring behaviour on the basis of object interactions in Section 2.4. In Section 2.5 we will describe the artificial intelligence techniques used in intelligent monitoring systems.

2.1 Requirements

A general architecture of an intelligent monitoring system is presented in Figure 2.1. The input of an intelligent monitoring system consists of the data from the sensors. Depending on the domain, specific criteria can be demanded of the sensors. For example, for some domains it is important that the sensors are cheap, very robust, or small.
The input data of the system sometimes needs preprocessing, because there could be errors in the data, or because not all data are relevant to the system. This preprocessing can be done on the sensor level, or on the level of the processing unit.

The processing part of the system is the intelligent part of the system and handles the data from the input. Domain knowledge is sometimes needed to reason about the data, for example for medical diagnoses (Milenkovic et al., 2006; Gao et al., 2005).
Artificial intelligence techniques are used for handling the data, for classifying the data, or

making decisions when and what to report, or about the actions to take. More about these techniques is described in Section 2.5 .

The output of the system consists of an interface for reporting, and/or actuators by which the system can take actions.
The system can have its own interface, or it is capable of connecting to other machines such as PDAs, PCs, or mobile phones through a wired or wireless connection. When using a wireless connection in combination with confidential information, then the security issues of such a wireless connection should also be taken into account.
When the system must be capable of taking actions, it needs actuators, such as for example valves that an irrigation monitoring system can open and close remotely (Damas, Prados, Gómez, & Olivares, 2001).

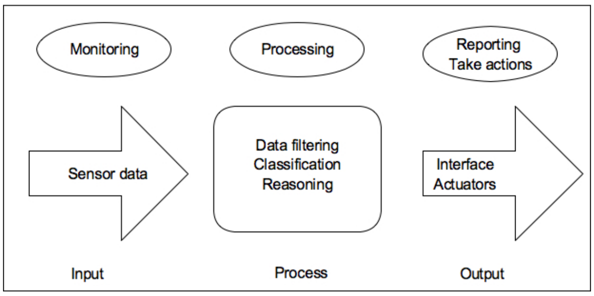

Figure 2.1: The input, process and output flow of the an intelligent monitoring system

The process of gathering data, processing this data, and taking actions occurs in a loop. The actions undertaken can have consequences on the input of the system. In some systems the input is constantly changing, independent of the output of the system, such as in the case of health monitoring. While the system is processing its data, the data can change. This has consequences on the actions that the system should take. A health monitoring system must be capable of handling these dynamically changing situations.

Depending on the domain, specific requirements are put on the system. For example when detecting the exact origin of a forest fire by using a sensor network (Akyildiz et al., 2002), the whole system must be very robust, and must be able to keep on functioning even when some sensors fail. When sensors fail, for example, because their view is blocked by a tree, the whole system must still function. This can be done because there are enough sensors that can take over the task of the blocked sensor.
Although reliability is important for all domains, in some domains it is extremely important, such as when monitoring the vital signs of patients (Milenkovic et al., 2006; Gao et al., 2005). The patient and the doctor must be able to rely on the correct functioning of the system.
For other domains it is important that the system is unobtrusive. For example in the domain of animal tracking (Nagl et al., 2003) where the system must not disturb the animals while monitoring them.

In some systems the processor of the system is a part of the sensor. These kind of sensors are called smart sensors. These sensors analyse the gathered data and take decisions based on that analysis. They may even be capable of monitoring their own condition; if they have enough power and if there are no errors in the measurements of the data. They can calibrate and optimize themselves (INCAS[3], 2009).

2.2 Sensors used

As already stated in Chapter 1 sensors can be categorized by the kind of quantity measured such as mechanical quantities, thermal quantities and chemical quantities (Harsányi et al., 2009). A diversion can also be made into wired and wireless sensors. Both kind of sensors have advantages and disadvantages. The advantage of wired sensors is that there is always a power supply possible, which is the biggest problem with wireless sensors. The advantage of wireless sensors is that they can be placed almost everywhere, no computer system or artificial device in the immediate vicinity is needed. The data can be transmitted through wireless communication such as bluetooth.
Depending on the domain a choice should be made for either wired or wireless sensors. The domain can also put extra demands on sensors, domains such as automatic irrigations (Damas et al., 2001) and monitoring forest fires (Akyildiz et al., 2002)) use hundreds or thousands of sensors to form a network covering a large area and therefore the sensors must be cheap.

2.3 Existing monitoring systems

In this section examples of intelligent monitoring systems will be described for different domains.

2.3.1 Environmental monitoring

In the field of environmental sciences measuring environmental parameters such as the water quality has been mostly done by pencil and paper notebook. This method is very labour-intensive and error prone. Vivoni and Camilli (2003) describe a system that streamlines the process of data collection. Sensors attached to mobile devices acquire the data from the environment. This data, together with the GPS location of the data, is communicated to a central data server for storage. In this central database the data from the different devices is integrated. In this way data collection, storages and merging can be done much faster and is less error prone.

2.3.2 Precision agriculture

Damas et al. (2001) describe an automatic irrigation system. In this system 1500 ha of land in Spain is monitored by 1850 control points which measure the hydratation of the ground. A control system supervises these control points in real-time and can open and close valves remotely, this way optimising the use of water and lowering the costs of exploitation.
Sensor systems are also used to assist the vineyard managers in making decisions about when to harvest the grapes, or when to spray with pesticides (Burrell, Brooke, & Beckwith, 2004).

2.3.3 Machine and process control

For plant operators it is often difficult to work efficiently with the many process values that are generated by computerised control systems. Computerised control systems are used to monitor, control and diagnose the different process variables such as temperature and pressure. Plant operators are faced with the task of monitoring these process data, analysing the current state, detect and diagnose process anomalies, and take appropriate actions to control the different processes.
In large-scale processes these tasks are difficult to perform efficiently, because of the large amount of data. An intelligent decision support system could be used to assist the plant operator with these tasks. The system will analyze these data and presents these data in a manner that reflects the important underlying trends or events in the process, thus making the tasks of the plant operator more manageable (Uraikul, Chan, & Tontiwachwuthikul, 2007).

2.3.4 Military applications

In military applications intelligent sensor systems are used for a variety of applications such as: monitoring friendly forces, equipment and ammunition; battlefield surveillance; reconnaissance of opposing forces and terrain; targeting; battle damage assessment and nuclear, biological and chemical attack detection and reconnaissance.
Sensor networks are, for example, used to monitor the status of friendly troops and the availability of the equipment and the ammunition. Small sensors attached to these parts give regular reports on the status of these parts.
Sensor networks are also used, for example, to cover a critical road with small sensors, in order to monitor activities of opposing forces on that road (Akyildiz et al., 2002).

2.3.5 Animal identification and health monitoring

Nagl et al. (2003) describe a wearable system that can be used to monitor the location and health of cattle. Sensors are used to measure parameters such as temperature, respiration and blood oxygen levels. Thereby monitoring the health of the individual animals. GPS is used to keep track of the location of the animals. Such a system is useful to be able to rapidly treat sick animals and prevent spread within or between herds.
A GPS system in the collar of a cow is also used by Butler, Corke, Peterson, and Rus (2004) to keep track of the animals and to keep them within a virtual fence. The collar consists of a sound amplifier, and when the cow reaches the virtual fence a sound is played. The sound will make the cow move away from the virtual fence. This way a herd can remotely be kept on one place, or slowly moved to another place.

2.3.6 Health applications

Baker et al. (2007) describe a system to prevent Sudden Infant Death Syndrome (SIDS), which can be caused because the child is sleeping on their stomach. The SleepSafe system described by Baker detects if the infant is sleeping on its back, side or stomach. This is performed by an accelerometer attached to the childs clothing. When the system detects that the infant is laying on its stomach or side, a warning signal is transmitted to the parents. This way the parents are offered peace of mind without feeling the urge to constantly check the infant while it sleeps.

To monitor vital signs of patients, wireless personal networks of sensors are used (Milenkovic et al., 2006; Gao et al., 2005). These systems monitor patients at home who are, for example, recovering from surgery, or are chronically ill. For most patients it is more comfortable to recover in their own home environment, and it is also less expensive then to admit them to a hospital. By monitoring the patients on a regular basis, the system is capable of detecting problems in the healthcare of the patients before the symptoms get worse. The system could also warn medical personnel when life-threatening situations occur. The system could provide patients with feedback on their health status and help to maintain an optimal health status. In Lorincz et al. (2004) a system is described which combines vital sign sensors with location tracking tags and a handheld computer in order to provide information, monitor the health status of patients and keep track of the location of the patients in case of a disaster. In this way multiple patients can be monitored and tracked at the same time.

2.3.7 Home applications

Intelligent sensor systems are also used frequently in home applications. A term frequently used is 'smart homes'. These are homes that are equipped with technology to improve the comfort of the residents, provide leisure and monitor the safety of the residents (Chan, Esteve, Escriba, & Campo, 2008). Comfort can be provided by automatically adjusting the lights when entering or leaving a room, turning down the radio when the phone rings, opening or closing curtains to adjust the light to the preferences of the resident (Soucek, Russ, & Tamarit, 2000).
Leisure can be provided by automatically recording television programmes that are interesting for the residents.
Safety is provided by monitoring the residents health, as described in the previous section about intelligent health monitoring systems.
Smart homes make it possible for the residents to receive continuous care at home, and therefore reduce the costs of placing them outside their own home, improving quality of life and enabling independence (Logan, Healey, Philipose, Tapia, & Intille, 2007).
Most smart home research focusses on the elderly people. The activities of daily living (ADLs) are monitored in order to asses an older adult's ability to remain independent (Dalal, Alwan, Seifrafi, Kell, & Brown, 2005). ADLs include eating, getting out of bed, using the toilet, bathing or showering, dressing, using the telephone, shopping, preparing meals, housekeeping, doing laundry, and managing medications (Tapia, Intille, & Larson, 2004). Most research focusses on detecting ADLs that are rather course grained. Focus on more fine-grained activities has been done by Patterson, Fox, Kautz, and Philipose (2005), who not only tried to recognizing that a meal has been prepared, but also what meal has been prepared. They centered on morning routines and included activities that are normally interleaved and share objects such as the knifes used for setting up the table and eating breakfast.

2.4 Monitoring behaviour through object interaction

The domain of monitoring the behaviour of people asks for an unobtrusive way of monitoring. People do not like the feeling that their privacy is affected and sensors must not hold them back in conducting the activities that they normally would perform. Cameras can feel obtrusive, but the most important reason why cameras are not suitable is that it is very hard to recognize on the basis of camera images what actions people are conducting, especially when the activities are recorded from different angles. Therefore another method is used, which is based on the 'invisible human hypothesis'. This hypothesis states that activities are characterized by the objects that are used during the performance of the activity (Patterson et al., 2005). Without seeing the human perform the action, the action can be recognised on the basis of the objects that are used by that human. The objects in the home are equipped with sensors to detect when these objects are used.

Monitoring object interactions can be done by Radio Frequency Identity (RFID) tags placed on the objects. Together with a RFID reader in a glove (Huang et al., 2008; Patterson et al., 2005; Philipose et al., 2004) or wrist worn bracelet (Logan et al., 2007) it can be detected which objects have been touched.
Another method is to use motion sensors on the objects to detect if the objects have been moved. When objects have been moved, it can be concluded that the objects have been used. The use of objects can also be measured by sensors that, for example, measure current flow to detect when electrical devices are used (Logan et al., 2007). The presence of people in the home can be detected by motion detectors, pressure mats and break beam sensors in the different areas of the home (Wilson & Atkeson, 2005).

The residents should not be hindered by the sensors and no extra effort of the user should be asked, therefore we will not use RFID tags in this thesis. We will use sensors in the objects to detect activities.

2.5 AI techniques for intelligent sensor systems

For the processing capabilities of intelligent monitoring systems Artificial Intelligence (AI) techniques are used. Input data need to be filtered, the system must be capable of reasoning about the data, or data need to be classified as belonging to a certain class (see Section 2.5.1).

A method that can be used for reasoning is an expert system (Argandoña et al., 2008). An expert system consists of a domain knowledge base that is structured and formalized in a way that the system can use it for reasoning. A common way to build such a knowledge base is by asking human domain experts about their knowledge and rules that they use when solving problems in the domain. This information is then formalized and stored in the knowledge base of the system.

Most intelligent monitoring system must be capable of handling uncertainty, because there can be errors in the data, or because not all information is available for the system. To handle uncertainty, techniques such as fuzzy logic or Bayesian networks are used. Fuzzy logic is a form of logic where the logical variables can take membership values of any value between 0 and 1, this way expressing the degree of certainty about a statement.
In Bayesian networks prior probabilities about events, such as that a fire starts, are combined with new data. The probability that an event has happened is calculated on the basis of the prior probabilities and the new information.

Because some domains, such as home automation, can consist of very large scale problems, multi agent systems are used to divide these problems into multiple small tasks. Different so called 'agents' in a multi agent system can be made responsible for these different tasks. Cook, Youngblood, and Das (2006) use multiple classification algorithms (the agents) and combine these algorithms in a multi-agent system for controlling the environment. The different agents are forced to work together to accomplish the overal goals of the system (Cook et al., 2006).

2.5.1 Classification methods

To classify which activities are being conducted different methods are used, for example a rule-based system (Dalal et al., 2005; Alwan et al., 2005). Rules are made about (combined) sensor values and the underlying activities based on observations of the target subjects and/or interviews with experts.
Another frequently used method is a learning algorithm that will construct a classifier based on a set of training examples with their class labels. A training example E consist of a set of attribute values $(x_1, x_2, ..., x_n)$ where x_i is the value of attribute X_i, and its class c. The goal of these learning algorithms is to learn to classify the class value c from the train data, and then to be able to make new classifications for unseen new data. A classifier should assign a class value c from the possible classes C to each new example.
The classes in the domain of activity recognition are the different activities. The input for these classifiers is the data from the sensors.
Classification algorithms that are used for classifying activities are for example decision trees (Logan et al., 2007; Bao & Intille, 2004; Isoda, Kurakake, & Nakano, 2004), naive Bayesian classifiers (Logan et al., 2007; Tapia et al., 2004; Bao & Intille, 2004) and support vector machines (Velera & Velastin, 2005).

Decision trees

A decision tree is called that way because it resembles the structure of a tree. The classifier starts at the root of the tree and will classify an example on the basis of performing a sequence of tests. After each test a branch of the tree will be followed until a leaf node has been reached. Each node in the tree corresponds to a test of the values of one of the attributes X_i. The branches from the node are labelled with the possible values of the test. Each leave node in the tree specifies a class value c.
For example in Figure 2.2 the nodes are “Mother present in kitchen”, and “Time bottle used”. The class values of the leaf nodes are “No” representing that no bottle has been prepared and “Yes” representing that a bottle has been prepared. The first test is if the mother is present in the kitchen. When this is not the case, the leaf node with the class value “No” is reached. When the mother is present in the kitchen the node wil be reached with the test of the time the bottle has been used. When the bottle has been used for less than 10 seconds, the leaf node with class attribute “No” will be reached. When the time the bottle has been used is longer than 10 seconds, the leaf node with the class attribute “Yes” is reached.

Decision trees are constructed by deciding which attribute has the highest information gain. A perfect attribute divides all the examples in the node into sets that contain only examples of one class. In the case of binary classification the class attributes have binary values such as for example“Yes” or “No”, 1 or 0, “True” or “False”. The algorithms for binary classification works as follows:

- If there are positive and negative examples, the best attribute to split these examples is chosen

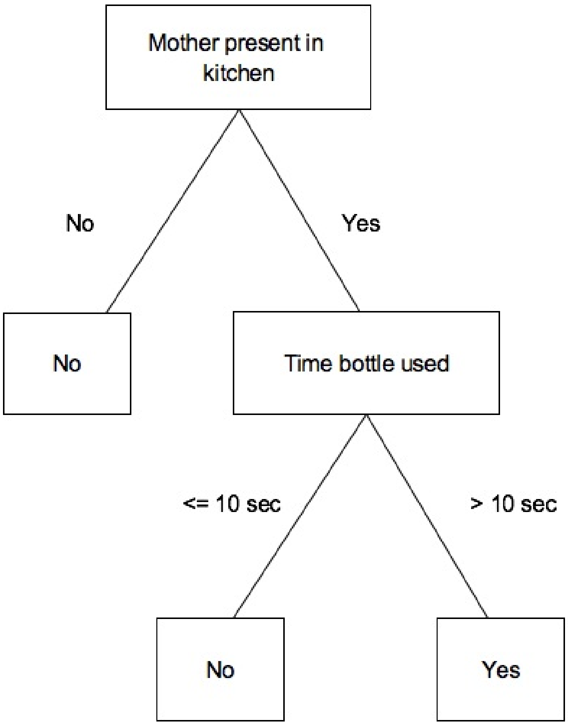

Figure 2.2: A decision tree for deciding if a bottle has been prepared

- If all the examples that are present in the node are positive, or all are negative, no further branching is needed. The node becomes a leaf node with the value 'yes' or 'no' as the decision for that leaf
- If there are no examples left it means that there was no example in the train set with those values. The best option is then to return the default value which is the majority classification calculated from the node's parent
- If there are no attributes left and still positive and negative examples, this can mean that some errors have occurred in the data, or that the attributes do not give enough information to describe the situation fully. The simplest way to solve this is to use majority vote calculated from the node.

The same kind of algorithm can also be used for multi class classification.
Logan et al. (2007) use decision trees for classifying activities conducted in a real-home environment. The data from the sensors is first converted into vectors, formed by concatenating all the data observed in a 30 seconds window, and with an overlap for each window of 15 seconds with the previous window. They use binary classifiers for each of the activities and assign each feature vector to Class 0 or Class 1 depending if the activity of interest had occurred at any time during the 30 seconds window covered. The WEKA software package (Weka, 2009) was used for experimenting with a decision tree and a naive Bayesian classifier. The results for the decision tree outperformed those of the Bayesian classifier.

Naive Bayes

A naive Bayesian classifier is based on Bayes' Rule. The probability that an example $E = (x_1, x_2, ..., x_n)$ belongs to class c is:

$$p(c|E) = \frac{p(E|c)p(c)}{p(E)}$$

When the class values only take the binary values of $+$ and $-$ it can be stated that E is classified as the class $C = +$ if and only if

$$f_b(E) = \frac{p(C = +|E)}{p(C = -|E)} \geq 1$$

The function f_b is called a Bayesian classifier (Zhang, 2004).
The classifier is called 'naive' when it assumes that the attributes are conditionally independent of each other, given the class. This assumption is rarely true in most real-world applications. In a naive Bayes network each attribute has only the class node as its parent, see Figure 2.3 for an example (Zhang, 2004; Russell & Norvig, 2003). The advantage of a naive Bayes classifier is that it can handle noisy data, and that the probabilities of the predictions can be given, which can be used as a report for the end user.

Tapia et al. (2004) have trained a naive Bayesian classifier to predict the activities conducted by residents in the residents own homes. They represented the data in a vector, which not only contained which sensors had fired during a certain time frame, but also temporal information, such as if a sensors had fired before another sensor. They used for each activity a different time window which has the length of the average time that activity lasted.
The results showed that the classifier scored the highest accuracy on the activities with the highest number of examples.

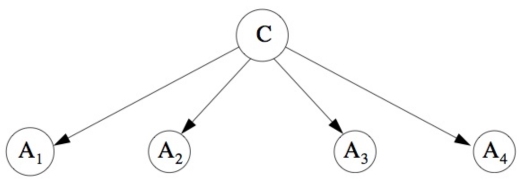

Figure 2.3: An example of naive Bayes (Zhang, 2004)

Support Vector Machines

The idea of a support vector machine (SVM) learning algorithm is to find the optimal hyperplane that separates two classes such that this hyperplane can be used as a decision boundary between those two classes. For example in Figure 2.4 the black circles are the negative examples and the white circles are the positive examples. The heavy line separates the black circles from the white circles. A line is sought that optimizes the distance of the line with the nearest circles of the negative and positive examples. These nearest positive and negative examples are called the support vectors.

In a lot of cases the data is not linearly separable in 2-dimensional space, but can be separated in a multi dimensional space. For example in Figure 2.5 it can be seen that the data can be separated by a circle, and in Figure 2.6 it can be seen that in a 3-dimensional space the data is linearly separable. Kernel functions are used to map the data points to a multidimensional space where the data points can be linearly separated. A good tradeoff is sought between the optimal hyperplane and the dimensions of the this new space (Russell & Norvig, 2003).

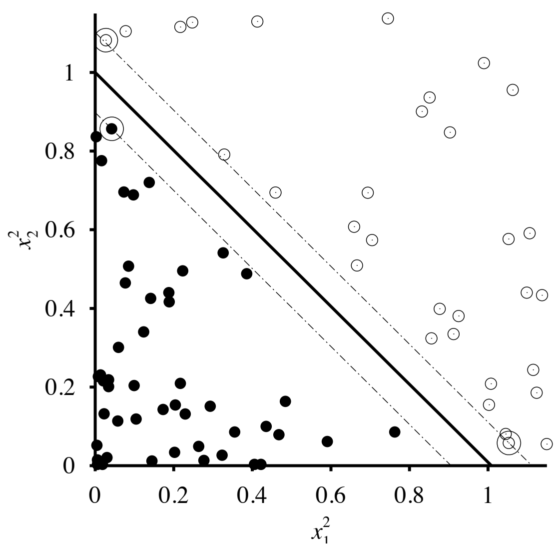

Figure 2.4: Finding an optimal linear separator that separates the positive (white circles) and negative (black circles) from each other (Russell & Norvig, 2003)

Huỳnh, Blanke, and Schiele (2007) use Support Vector Machine algorithm to classify activities. The input for the classifier are the histograms of sensor values over a period of time. The results showed that the SVM algorithm outperformed the other classifiers used.

2.5.2 AI approaches used in this thesis

One of the research questions to be answered in this thesis is how reliable can it be classified that a bottle has been prepared (RQ2). Therefore we will experiment with different classifiers to test how well they are able to classify the tasks from the data of the experiments. As described in the previous section (Section 2.5.1) decision trees, naive Bayesian classifiers and support vector machines have shown good results on classifying activities. We will therefore test these three classifiers in this thesis. More about the way we used these algorithms is described in Chapter 4.

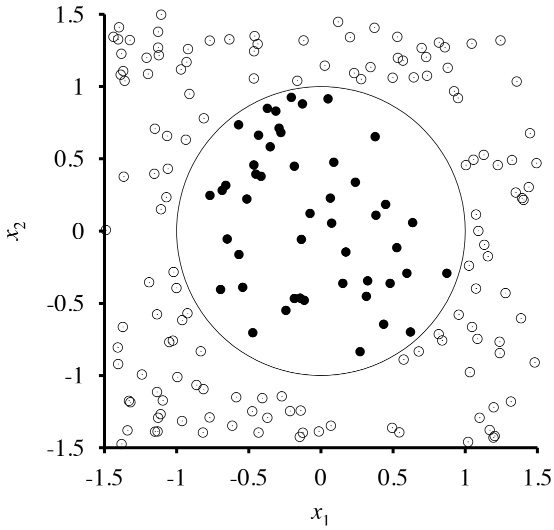

Figure 2.5: Data that is not linearly separable in the 2-dimensional space (Russell & Norvig, 2003)

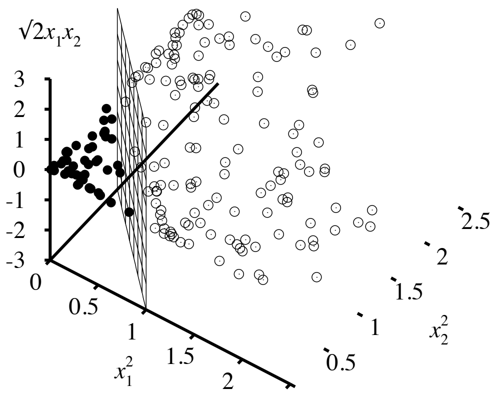

Figure 2.6: Same data as in Figure 2.5 is linearly separable in the 3-dimensional space (Russell & Norvig, 2003)

Chapter 3

Research environment and experimental design

The research of this thesis focuses on an intelligent monitoring system to support people with a mental disability. The problems of these people are very diverse. We have therefore chosen to center on one problem group, the group of mentally disabled mothers. These mothers live on their own and need support in their daily life and with the upbringing of their child. Problems in this group occur in the areas of hygiene (giving the child a new nappy on time, wash the child), discipline (mother does not now how to correct the child), general attention for the child (child is left alone for a long time) and food (gives only unhealthy food to the child or forgets to give the child something to drink) (Meulen, Schaaf, Tan, & Wegman, 2008) (J. van der Burgh, teamcoordinator NOVO ambulant team, personal communication, October 8, 2008).

In Section 2.4 we described the research into detecting activities in the home on the basis of objects and why this method should be used for monitoring behaviour. We have chosen to focus on the problem of monitoring that the mother gives the child something to drink, because that task seemed to be well defined with respect to the objects used during that task.

In order to classify the task of giving the child something to drink we needed a relevant dataset. No relevant data set was available, therefore we have created a dataset ourselves. We built a test environment and asked several test subjects to perform tasks in this test environment. We first had to define the tasks that the subjects had to perform, which we will describe in Section 3.1. Based on this task analysis we could built our test environment and equip the required objects with sensors. This is described in Section 3.2.
In Section 3.3 we will describe the software platform built for the data acquisition, in Section 3.4 we will describe the experiments conducted.

3.1 Task analysis

We first made an analysis of the actions and the objects that play a role in giving a child something to drink. A diagram of this analysis can be seen in Figure 3.1. The actions that needed to be conducted can be divided into; fetching the drink, making the drink and the child drinking. Drinks can be fetched from the refrigerator, from somewhere outside the refrigerator or water can be tapped from the water tap. For the preparation of a drink the bottleheater, microwave, electric kettle or cooker can be used, and a cup, bottle or glass can be fetched. When the child drinks the bottle is near the child, the liquid level in the bottle decreases, and the angle of the bottle can indicate drinking (a slanting bottle indicates drinking, straight indicates no drinking). An indication can also be that the child holds the bottle, or that the bottle/cup is held to the mouth. The use of the child seat or the feeder can indicate that the child gets something to eat or drink. The presence in the kitchen can suggest that the mother is preparing food or drinks. The child crying can inform that the child is hungry or thirsty.

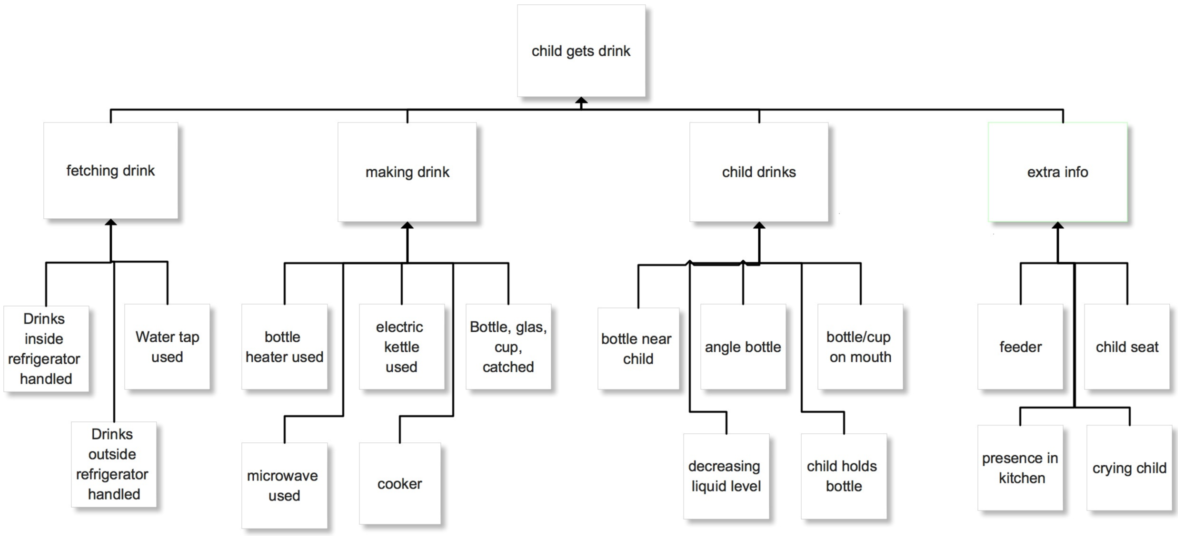

Figure 3.1: Analysis of the actions and objects used in giving a child something to drink

As can be observed from Figure 3.1 there are still many possible objects that could be used and actions that could be undertaken when providing a child with something to drink. The problem domain is too broad for the scope of this study, we therefore decided to narrow down this problem domain.

We decided to focus in this research on preparing a bottle for a child. This way only a limited amount of objects would be needed and we could define a small task set for our test subjects. As it turned out, there are different ways of preparing a bottle (Thebe, 2009). In general a bottle must be prepared by mixing powder milk with water in the bottle and warm the bottle. One way is to mix the powdered milk with cold water and then warm the bottle. A second way is to first warm the bottle with water and then add the powdered milk. See Figure 3.2 for a schematic overview of the ways of preparing a bottle. Based on these different ways of preparing a bottle we defined the tasks of the test subjects, which is described in the next section (Section 3.1).

3.1.1 Tasks for the test subjects

The tasks that the test subjects had to perform contained the four different ways of warming a bottle:

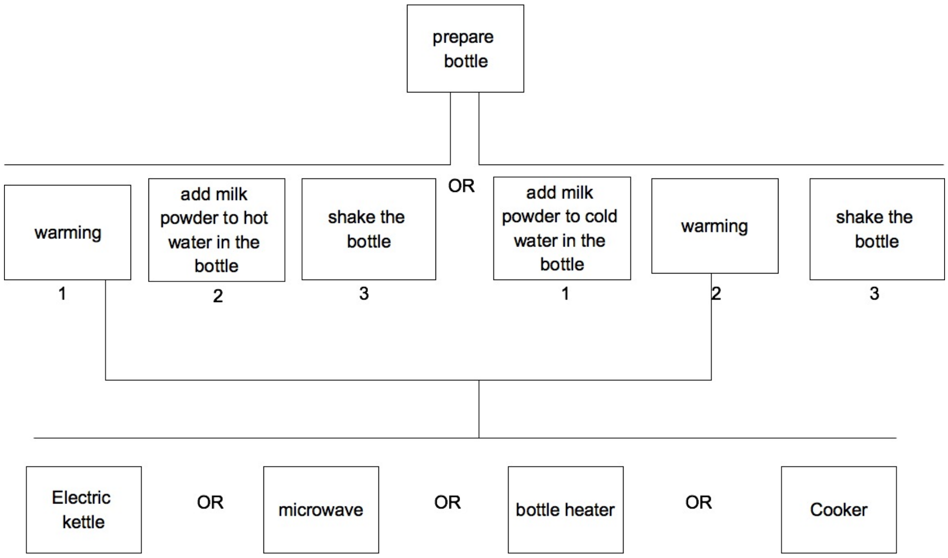

Figure 3.2: The different ways of preparing a bottle for a child

- Use an electric kettle to boil the water first and use this water for preparing a bottle
- Use the bottle heater for warming the bottle
- Use the microwave for warming the bottle
- Warm the bottle in a pan with water on the cooker

In order to test if the system was capable of distinguishing the bottle making tasks from other tasks performed in the kitchen, some other (not related) tasks have been added to the experiment. There are of course a lot of tasks that can be performed in the kitchen, so it has been decided to choose those tasks that have a certain kind of overlap with the bottle making tasks. The idea is that if the system is capable of detecting bottle making tasks between these other partly overlapping tasks, the system is certainly capable of detecting bottle making tasks amongst less overlapping tasks. Because the tasks will be recognized on the basis of object use, the overlapping tasks are chosen on the amount of similar objects that are used. For example preparing tea will trigger the sensor of the kettle and the sensors in the kitchen cabinets where the tea and cups are stored. The same sensors will be triggered when preparing a bottle by using the kettle to boil the water, but in the last case, also the sensor of the bottle will fire. Six tasks that all share one, two or three objects with the tasks of bottle making are listed below:

- Prepare tea with the kettle
- Warm chocolate milk in the microwave
- Warm soup on the cooker
- Clean the bottle and put the bottle back in kitchen cabinet
- Clean the cups and put them back in kitchen cabinet

- Grap a cup from the kitchen cabinet and fetch some water from the water tap

This way we constructed 10 tasks that the subjects were supposed to perform in the kitchen, see Table 3.1 for an overview of these tasks.

Table 3.1: Overview of the tasks for the subjects

Task name	Description
bottleBottleheater	prepare a bottle and use the bottle heater for warming the bottle
bottleCooker	prepare a bottle and use the cooker for warming the bottle
bottleMicrowave	prepare a bottle and use the microwave for warming the bottle
bottleWaterboiler	prepare a bottle and use the electric kettle for warming the bottle
chocolatemilkMicrowave	warm chocolate milk in the microwave
cleanBottle	clean the bottle and put the bottle back in kitchen cabinet
cupWater	take a cup from the kitchen cabinet and some water from the water tap
dishWashing	clean the cups and put them back in kitchen cabinet
soupCooker	warm soup on the cooker
tea	prepare tea with the waterboiler

3.2 The objects and sensors used in detecting the preparation of a bottle

Based on the task analysis described in Section 3.1 we defined the following objects that played a role: a bottle, electric kettle, bottle heater, microwave, cooker and pan. Also kitchen cabinets are used to store the bottle and the powdered milk, and the watertap is used for water. In order to measure that these objects are used, the following sensors were chosen:

- *Electronical devices:* We decided to use ammeters (see Figure 3.4) to measure the electric current of these devices. The electric current gives the best information about if the device is turned on.

- *Doors of the kitchen cabinets:* To monitor if the doors were opened or closed, we used magnet contact sensors (see Figure 3.5). We placed magnet contact sensors in the kitchen cabinets and a magnet on the doors. The magnetic radiation measured by the sensor gives a precise indication if the doors are closed or opened. We used a threshold value for the sensor to make a distinction between open and closed.

- *Water taps:* To measure the opening and closing of the water taps we used magnet contact sensors. The magnet contact sensors were placed under the hot and cold water tap and a little magnet was placed on both watertaps (see Figure 3.6). When the tap twisted, the magnet contact sensor would sense a differences in the magnetic field. We first tried to use a vibration sensor on the watertap to measure when the tap was twisted, but this did not work, because the rotation of the tap did not cause enough movement for the vibration sensor to detect. The solution with the magnets worked fine, even when the tap was twisted only a small amount, the sensor sensed this change in radiation. We used a threshold value to make a distinction between if the taps were opened or closed. When the tap was twisted too much the value of the sensor would decrease below the threshold value, wrongfully indicating that the tap was closed, but because the sensor value would first reach the threshold value before dropping below this value, it still indicated that the tap was opened.

- *Cooker:* A rotation sensor was used for the cooker (see Figure 3.7 and 3.8) to detect when the cooker was switched on. We also used a threshold value for this sensor.

- *Bottle:* A vibration sensor was used to detect if the bottle was used (see Figure 3.9). The vibration sensor was placed on the top of the bottle. This top can be removed from the bottle, making it possible to place the bottle in the microwave without the sensor. This was necessarily because electronic and iron parts should not be placed in a microwave. The test subjects in the experiment are forced to put the top back on the bottle before they can shake the bottle without spilling milk, making sure that the sensor of the bottle will be triggered both at the beginning and the end of the task. Because the wire of the sensor was not long enough, it was not possible for the test subjects to put the bottle in the microwave with the sensor still on it.

- *Motion:* We used a motion detection sensor (see Figure 3.10) to detect the presence of the subjects in the kitchen. Because the subjects were always present in the kitchen during the tests, this information turned out to be not relevant in our experiments.

For the sensors it has been decided to use phidgets (Phidgets, Inc., 2008). Phidgets are a set of ‘plug and play‘ building blocks for low cost USB sensing and control by a PC. They come with a lot of documentation and can be connected through an interfacekit (see Figure 3.11) that takes care of the communication with the PC.

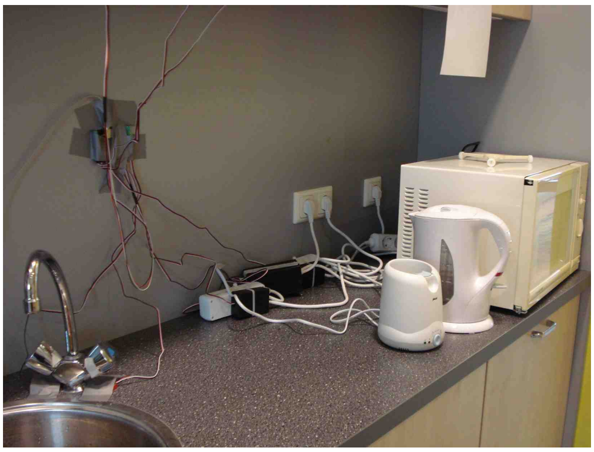

Figure 3.3: Overview of the kitchen environment

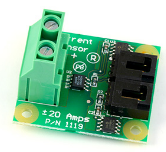

Figure 3.4: Ammeter (Phidgets, Inc., 20Amp Current sensor)

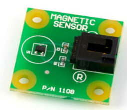

Figure 3.5: Magnet contact sensor (Phidgets, Inc., linear Hall Effect Sensor)

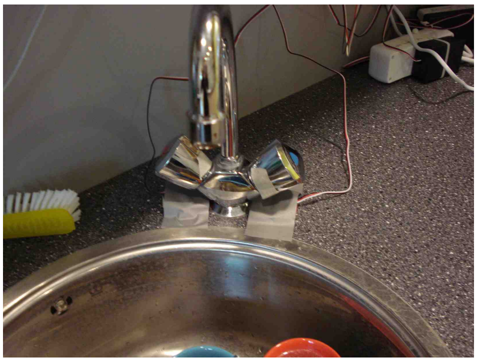

Figure 3.6: Magnet contact sensors used to measure twisting of the watertaps

Figure 3.7: Rotation sensor (Phidgets, Inc.)

Figure 3.8: Rotation sensor attached to the interfacekit and the cooker

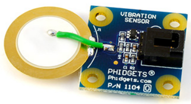

Figure 3.9: Vibration sensor (Phidgets, Inc., Piezoelectric Vibration sensor)

Figure 3.10: Motion sensor (Phidgets, Inc., Infrared motion sensor)

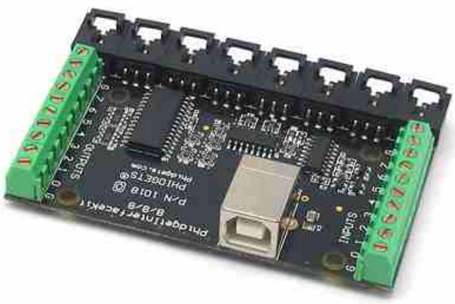

Figure 3.11: Interfacekit (Phidgets, Inc.)

3.3 Data acquisition process

In order to store and label the data from the experiment, a test platform was built. The sensors were connected to a computer that runs this test platform. The platform takes care of reading in the sensor data. The interfacekit automatically gave a signal when one of the sensor values changed. When the vibration sensor of the bottle or the motion sensor fired, this meant that this object was used. For the other sensors the value was compared to a threshold value to decide if the object was opened (kitchen cabinets, watertaps) or used (kettle, bottle heater, cooker, microwave). The test platform also had a user interface, on which the experimenter could monitor if the sensors were connected and were working correctly. The user interface also consisted of check boxes for the ten different tasks, on which the experimenter could check what task was being performed.

3.4 The experiments conducted

The experiments have been conducted in the test kitchen. We have asked several test subjects to perform the tasks described in Section 3.1 in this kitchen. In a real world setting most people will perform multiple tasks at the same time in a kitchen. We therefore conducted, besides a round where the tasks where supposed to be performed on-by-one, a round where the subjects were allowed to perform multiple tasks at the same time. This way we gathered a data set containing single task data from the first round, and a data set of multi task data from the second round. The two rounds were performed after each other with some minutes break between the rounds. During this break the objects in the kitchen were put back on their place if required.

3.4.1 The test environment

A kitchen at CENSI has been used for the experiment. This kitchen consisted of multiple kitchen cabinets, a refrigerator and a kitchen sink. For the purpose of the experiments also a cooker, microwave, electric kettle and bottle heater were added to the environment of the kitchen. Only two kitchen cabinets have been used for the experiment, one kitchen cabinet was used for the cups and the other was used for the tea and milk powder.

3.4.2 The test subjects

Eleven subjects have participated, five females and six males. The ages ranged from 19 to 56 with a mean of 27. The subjects were recruited from the direct colleagues, students and teachers present in the building of CENSI. The student subjects were paid 6 euro on the completion of the experiment, which took about half an hour.

3.4.3 The experimental procedure

At the start of the experiment the subjects received written instructions (Appendix A). After reading the instructions the subjects had the opportunity to look around in the kitchen and were orally instructed where everything was situated and how the different devices would work. The experiment consisted of two rounds. During the first round the subjects were instructed with only one task at the time. During the second round the same tasks were

given to the subjects, but this time, multiple tasks were given at the same time.
The different tasks were printed on 10 cards, see Appendix C, and these cards were shuffled before the experiment. The experimenter takes the top card and gives this card to the subject. The experimenter checks the task on the test platform interface so that the data was labelled correctly. The subject performes the task and when he or she is ready, the card is given back to the experimenter and the experimenter unchecks the task in the interface. Then the next card is taken from the top, until all cards have been used. In the second round the cards are again shuffled and this time 3 or 4 cards at the same time are given to the subject. The subject is allowed, but not obligated, to perform multiple tasks at the same time.

Because most subjects were not familiar with preparing a bottle, instructions about this task were printed for the subjects, see Appendix B. There were two versions of the instructions, in one version the instructions involved first warming the water and then adding the powder milk. In the other version the instructions were to add the powdered milk first and then warm the botte.

3.4.4 The problems during the experiments

Some problems with the sensors and with the subjects arose during the experiment, these problems are described below.

Problems with the sensors

The problems with the sensors had to do with:

- *Sensor malfunctioning.* For example the sensor of the bottle heater kept indicating that the bottle heater was on, while it was turned off. We corrected this error manually in the data afterwards.

- *Constraints on the sensors because of the environment*: For example not all sensors that we wanted for the bottle could be used. Sensors can not be placed under water, or in a microwave, and those are actions the bottle would be exposed to. In our first set-up we wanted to use a temperature sensor to measure the temperature of the bottle. We also wanted to use a light sensor in combination with a led light, in order to measure the density of the liquid inside the bottle. This would indicate if the bottle contained milk. Because of the constraints that the bottle must be able to be placed in the microwave, and under water, it has been decided to only use a vibration sensor on the top of the bottle. This top could be removed when the bottle was placed inside the microwave or under water. Because the wire of the vibration sensor was not long enough, it was just not possible to place the bottle inside the microwave with the top still on it.

- *Problems caused by test subjects*: Problems occurred because some test persons held the bottle with sensor under water. This caused the vibration sensor to fail. We had to replace this sensor with a new vibration sensor. We used extra tape with this new sensor to protect the sensor from water.

Problems with the subjects

Some unexpected situations occurred during the experiment that we did not correct, because we thought these were normal variations in the way people conduct the different tasks and this would bring some variation in the data. These situations were:

- One subject indicated that he did not like hot chocolate and therefore used the microwave only for a very short moment to warm the chocolate milk.

- Three subjects used the pan to warm water and put the water from the pan in the bottle, instead of warming the bottle au bain-marie.

- One subject used the pan to heat the water for tea, instead of the electric kettle.

Chapter 4

Classification process and results

In this chapter the classification process and the results of the experiments will be described. The data flow of the classification process is presented in Figure 4.1. In Section 4.1 we will first explain how the data was filtered and segmented, followed by Section 4.2 where the representation of the data into vectors will be discussed. In Section 4.3 we will describe the normalization and standardization process. In Section 4.4 the WEKA (Weka, 2009) environment and the classification algorithms used will be described.
Followed by Section 4.6 where the results of the classification algorithms will be presented.

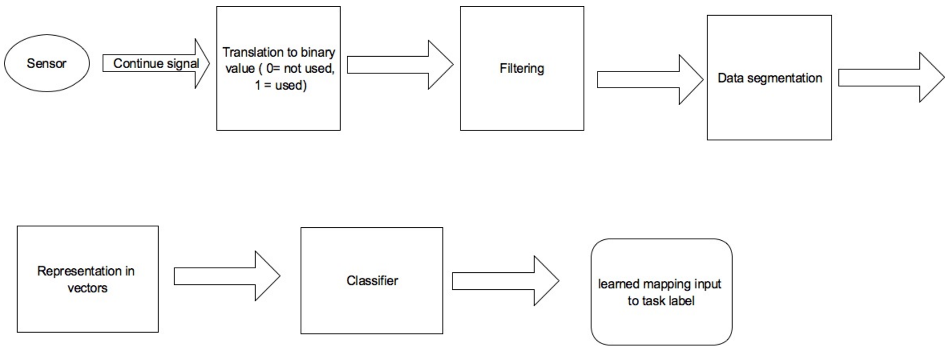

Figure 4.1: The data flow

4.1 The processing of the raw data

The output of the sensors consists of a continuous signal. In the test platform this continuous signal was translated into a binary signal, this because we were only interested in if the objects were used or not used. Examples of these binary signals for the different tasks are presented in Figures 4.2, 4.3 and 4.4. In Figure 4.2 it can be seen that first kitchen cabinet 1 is used, followed by bottle, water tap cold, microwave, kitchen cabinet 2, bottle and again kitchen cabinet 2. In Figure 4.3 it can be seen that first kitchen cabinet 2 is used for a short moment of time and then the cooker. In Figure 4.4 it can be seen that first the water tap cold is used and then the electric kettle (water boiler) is turned on. During the time the kettle is on, kitchen cabinet 1 and 2 are used.

The raw data consisted of the task label and for each sensor a binary value expressing if the sensor had fired during the past one second. These raw data was first edited by hand to filter out some errors in the data.
One of the errors that we had to filter out was that the sensor of the bottle heater kept firing although the bottle heater itself was switched off. This was caused because the value of the ammeter for the bottle heater was very low when the bottle heater was turned on. The current sensor was not able to distinguish this low value from the value when the bottle heater was turned off. We manually changed the ones to zeros in the data where this error occurred.

During the experiment the subjects were allowed to turn the bottle heater and cooker off, when they had nothing else to do than to wait for the minutes to pass that it would normally take for the bottle heater and the cooker to warm the bottle and the soup. We manually prolonged the time for these devices to the average time these devices would normally be used.

We labelled the task at the moment the subject received the instruction. A better definition of the start of a task would be the first sensor firing, because some time passed between the moment that the subject received the instructions and the moment the subject started with the task. Therefore the data was edited in such a way that the task started at the first sensor firing and ended at the last sensor firing during that task.

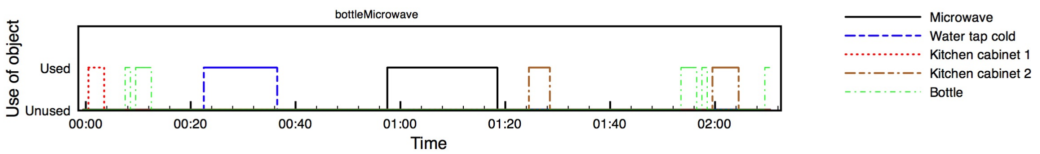

Figure 4.2: Object use during the task bottleMicrowave

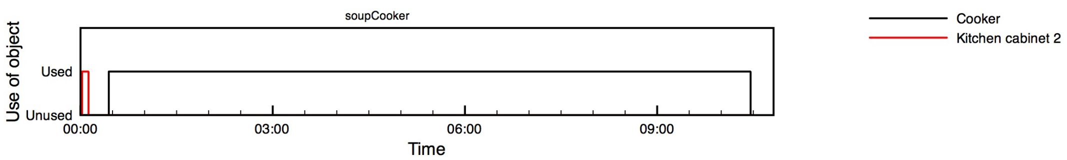

Figure 4.3: Object use during the task soupCooker

4.2 Representation of the data

The next step was to find a way of representing the data. As can be seen in Figures 4.2, 4.3 and 4.4, there is a sequence in time of objects that are used. A way to represent the data is to take this sequence into account.

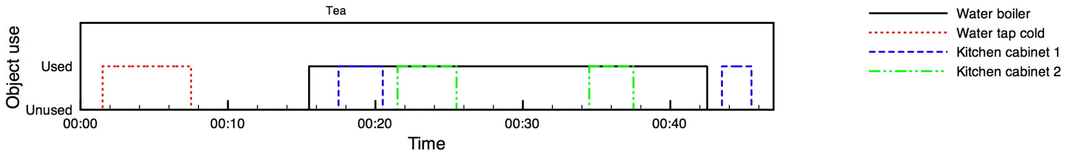

Figure 4.4: Object use during the task tea

We decided to first examine a more basic representation, and when needed, we could later on examine the more complex representation of sequences of objects in time.
Each task of each subject was being represented as a vector consisting of the label of the task, the total amount of time that the task has lasted and for each individual sensor the total time that that sensor had fired during the task. An example of such an representation into vectors can be found in Table 4.1. The left column are the task labels, the second column is the amount of time the task has lasted and the following columns are the amount of time the sensors of the different objects have fired.

Table 4.1: Example of the representation in vectors

	duration	microwave	waterboiler	bottleheater	cooker	watertap hot	watertap cold	kitchen cabinet 1	kitchen cabinet 2	bottle	motion
bottleBottleheater	562	0	0	478	0	0	38	7	3	23	529
bottleCooker	635	0	0	0	604	0	6	0	5	12	619
bottleMicrowave	107	33	1	0	0	0	8	0	5	15	84
bottleWaterboiler	91	0	19	0	0	0	7	0	11	16	68
chocolatemilkMicrowave	24	23	0	0	0	0	1	0	0	8	17
cleanBottle	22	1	0	0	0	0	5	0	0	13	20
cupWater	4	0	0	0	0	0	3	0	0	1	2
dishWashing	12	0	0	0	0	11	0	1	0	0	8
soupCooker	619	0	0	0	607	0	0	0	2	1	615
tea	100	0	99	0	0	1	0	4	11	3	47

4.3 Normalization

A possible next step is to 'normalize' or 'standardize' the vectors. 'Normalizing' a vector means rescaling the vector by the minimum and range of the vector, to make all the elements lie between 0 and 1. This is done by dividing the elements in the vector by the highest value.

$$e_i = \frac{v_i - \mathrm{min} v_i}{\mathrm{max} v_i - \mathrm{min} v_i}$$

Where v_i is the actual value of element i and the maximum and minimum are taken over all elements.

'Standardizing' a vector means rescaling the vector to have a mean of 0 and a standard deviation of 1. This is done by subtracting the mean from the elements in the vector and divide the elements by the standard deviation of the elements (Sarle, 2009; Witten & Frank, 2005).

$$e_i = \frac{v_i - \overline{v}}{\sigma}$$

Where v_i is the actual value of element i ,$\overline{v}$ is the mean of the elements and σ is the standard deviation of the elements.

4.4 The classification algorithms used

As discussed in Section 2.5.2 we decided to test three classifiers, a decision trees classifier, a naive Bayesian classifier and a support vector machine. These classifiers were implemented in WEKA, which is described below.

4.4.1 The WEKA environment

The WEKA environment is an open source software environment consisting of a collection of machine learning algorithms for data mining tasks (Weka, 2009). The algorithms we wanted to use are all implemented in this environment.
The data had to be translated to the ARFF data format required by WEKA.

Support vector machine

The support vector machine algorithm implemented in WEKA is based on the algorithm described by Platt (1998) and modified by Keerthi, Shevade, Bhattacharyya, and Murthy (2001) and is called Sequentially Minimalization Optimization (SMO). These modifications have to do with the speed of calculations and the ease of implementation of the algorithm.
The default parameter settings of WEKA for this algorithm have been used. This meant a complexity constant of 1, a tolerance parameter of $1.0e^{-3}$, an epsilon for round-off error of $1.0E^{-12}$ and the use of a polykernel with exponent 1.
The parameter that we varied was if the vectors should be normalized, standardized, or neither. We tested all these three options.

Decision tree

The J48 algorithm implemented by WEKA is an implementation of the C4.5 decision tree algorithm described by Quinlan (2003). We used the default settings for this algorithm; a confidence threshold for pruning of 0.25 and minimum number of instances per leaf of 2.

Naive Bayes

A naive Bayesian classifier is implemented by WEKA. A description of a naive Bayesian classifier is given in Section 2.5.1. There are no parameters for this function.

4.5 Method for evaluating the classifiers

Evaluation of the classifiers is done by 10-fold cross-validation. First the data is split into 10 folds. Training is done on 9 folds and the testing is done on the left out fold. This is repeated 10 times, each time testing on another fold.

We used a corrected resampled t-test to analyse if there was a significant difference in the results of the classifiers.
The corrected resampled t-test is used to circumvent the problem that when you use repeated cross-validation, the differences will eventually be judged significant, by just increasing the amount of runs. This because the value of the t-statistic will increase without bound. To prevent this a corrected t-test is used:

$$t = \frac{\overline{d}}{\sqrt{(\frac{1}{k} + \frac{n_2}{n_1})\sigma_d^2}}$$

Where n_1 instances are used for training, n_2 instances are used for testing. The variable $\overline{d}$ is the difference between the means of the percentages correctly classified instances for the classifiers that are compared. The variable k is the amount of runs. In the case of 10-fold cross-validation which is repeated 10 times, $k = 100$ (Witten & Frank, 2005).

4.6 Results

To see how well the tasks could be classified, we tested the three classification algorithms described in the previous Section 4.4 on the data from the experiments. The input for the classifiers are the vectors as described in Section 4.2.
First a decision had to be made about normalizing or standardizing the data, which is described in Section 4.6.1, followed by the presentation of the results for the single task classification in Section 4.6.2 and the results of the multi task classification in Section 4.6.3.

4.6.1 Normalizing or standardizing the data

We tested the option of normalizing and standardizing the data with the SVM algorithm. The features, such as the time the microwave has been used, are normalized/standardized. The results are presented in Table 4.2. We used the corrected paired t-test to analyse if there was a significant difference in the results of the three normalization options.

Standardization and no scaling score significantly better on the 0.01 significance level. There is no significant difference between the standardization an no scaling option on the 0.05 level. The reason that normalization scores significantly less can be because of the fact that all features are rescaled to lie within 0 and 1. This way information will be lost about striking high values of features for certain tasks, such as the high value of the time the cooker is used in the task of preparing soup.
Because there was no significant difference between the standardization option and no scaling, we have chosen to continue our analysis with the non-scaled data.

Table 4.2: Results of normalizing, standardizing or no scaling of the data

Method	Percentage Correctly Classified Instances
Normalization	71.8%
Standardization	94.2%
No scaling	94.1%

4.6.2 Single task classification

We first looked at the least complex challenge, the classification of the single tasks.
We trained the classifiers on all the single task data, and we used 10-fold cross-validation to test the classifiers. The results of this method are presented in Table 4.3.

Table 4.3: Results for 10-fold cross validation on all single task data

	SVM	C4.5 decision tree	Naive Bayes
Percentage Correctly Classified Instances	94.1%	93.6%	93.6%

We used a corrected paired t-test to analyse if there was a significant difference in the results of the three classifiers. There was no significant difference on the 0.05 level between the three classifiers. It has been chosen to use the SVM classifier for the rest of the analysis, this classifier has the highest percentage correctly classified instances.

An analysis has been made about which tasks were misclassified. The confusion matrix of the SVM is presented in Table 4.4. From this confusion matrix it can be concluded that bottleBottleheater, tea, cupWater, soupCooker and bottleCooker are never misclassified. Dishwashing is once misclassified as cupWater and once as cleanBottle. This makes sense because dishWashing, cupWater and cleanBottle all share the use of the water tap. Also the mis classification of cleanBottle with dishWashing makes sense, because they share the use of the watertap. Bottlewaterboiler is once misclassified as tea, which also makes sense because they both use the waterboiler. An explanation for the mis classification of bottleMicrowave as cleanBottle is that they both use the bottle.
No direct explanation can be given for the mis classification of chocolatemilkMicrowave as dishwashing.

We used also a second method of training and testing; we trained the classifier on all single task data except the data from one subject, and tested the classifier on the data of this left out subject. We repeated this leave-one-subject-out cross-validation for all subjects. The results presented in Table 4.5 show that the SVM classifier scores on average 94.5% correct with this method. These results also show that all subjects have been conducting the tasks in a similar way, because in 10 cases the classifier scored above 90% correct classified tasks while it was trained on the tasks performed by other subjects. It can also be seen that the score for subject 3 is the least with 80% correct. From the confusion matrix of this subject (Table 4.6) it can be seen that preparing chocolatemilk in the microwave was misclassified as doing dishwashing, but in the notes from the experiments it can be found that subject 3 had

Table 4.4: Confusion matrix for 10-fold cross-validation on all single task data for the SVM algorithm

a	b	c	d	e	f	g	h	i	j	<– classified as
11	0	0	0	0	0	0	0	0	0	a = bottleBottleheater
0	9	0	0	1	0	1	0	0	0	b = dishWashing
0	0	11	0	0	0	0	0	0	0	c = tea
0	0	0	10	0	0	1	0	0	0	d = bottleMicrowave
0	0	0	0	11	0	0	0	0	0	e = cupWater
0	1	0	0	0	10	0	0	0	0	f = chocolatemilkMicrowave
0	2	0	0	0	0	9	0	0	0	g = cleanBottle
0	0	0	0	0	0	0	11	0	0	h = soupCooker
0	0	1	0	0	0	0	0	10	0	i = bottleWaterboiler
0	0	0	0	0	0	0	0	0	11	j = bottleCooker

put the microwave on for a very short moment of time on purpose because he did not like hot chocolate milk. Other mis classifications have been discussed in the previous paragraph.

Table 4.5: Results single task classification for the SVM classifier with leave-one-subject-out cross-validation

Test person	percentage correctly classified
1	90%
2	100%
3	80%
4	90%
5	100%
6	100%
7	100%
8	100%
9	90%
10	100%
11	90%
average	94.5%

Table 4.6: Confusion matrix SVM classifier single task classification leave-one-subject-out cross-validation for subject 3

a	b	c	d	e	f	g	h	i	j	<– classified as
1	0	0	0	0	0	0	0	0	0	a = bottleBottleheater
0	1	0	0	0	0	0	0	0	0	b = dishWashing
0	0	1	0	0	0	0	0	0	0	c = tea
0	0	0	1	0	0	0	0	0	0	d = bottleMicrowave
0	0	0	0	1	0	0	0	0	0	e = cupWater
0	1	0	0	0	0	0	0	0	0	f = chocolatemilkMicrowave
0	1	0	0	0	0	0	0	0	0	g = cleanBottle
0	0	0	0	0	0	0	1	0	0	h = soupCooker
0	0	0	0	0	0	0	0	1	0	i = bottleWaterboiler
0	0	0	0	0	0	0	0	0	1	j = bottleCooker

Because the main task was to classify that a bottle has been made, two groups have been made, preparing a bottle, and other tasks. The confusion matrix for these groups is presented in Table 4.7. The accuracy measures are presented in Table 4.8.

Table 4.7: Confusion matrix bottle vs other tasks

bottle	other	<– classified as
42	2	bottle
0	66	other

Table 4.8: Accuracy measures bottle vs other tasks

True positive rate	False positive rate	class
0.95	0	bottle
1	0.05	other

From this table it can be seen that the true positive rate for bottle making is very high, in 95% of the cases that a bottle has being prepared, this has been recognized as such. The false positive rate is 0.

Conclusions

From the results it can be concluded that the single tasks can be recognized very well, all three classifiers score above the 93% correctly classified tasks.
From the confusion matrixes in Table 4.4 it can be concluded that most mis classifications occur between tasks that were very similar during the experiment (four of the six mis classifications). Other tasks were also similar, but were correctly classified by the SVM classifier.
When the tasks are grouped by preparing a bottle versus other tasks, the SVM classifier scores a false positive rate for the bottle making group of 0, so in 0% of the cases a task has been classified as preparing a bottle for a child, when this was not the case. This is very important, because in the eventual monitoring system the costs for false positives is high. In the case of false positives, the system will not remind the mother, when this is necessary, because the system will think that a bottle has been prepared. The costs for false negatives is not that high, the system will remind the mother that a bottle should be prepared, when she has already prepared a bottle.

4.6.3 Multi task classification

For the multi task classification we also used the SVM classifier. Multiple classifiers have been trained, for each task one classifier. An analysis has been made of which tasks were easy to recognize and which tasks were hard to recognize. We first trained only on the multi task data and tested the classifier by using 10-fold cross-validation on the multi task data, these results are presented in Table 4.9. We also trained on all data (single task data and the multi task data together) and used 10-fold cross-validation to test the classifier on all these data, for the results see Table 4.10.
It can be seen from Table 4.9 that the hardest task to recognize among the multi tasks was cleanBottle, with a recognition of 60%. This is also the hardest task to recognize when trained and tested on all the data. with a recognition percentage of 81%.

The task of recognizing the task bottleBottleheater is recognized in 100% of the cases for the both multi task data and for all the data.

Table 4.9: Results for 10-fold cross-validation on multi task data

Task	Correctly classified
bottteBottleheater	100%
dishWashing	76 %
tea	69 %
bottleMicrowave	67 %
cupWater	73 %
chocolatemilkMicrowave	80 %
cleanBottle	60 %
soupCooker	78 %
bottleWaterboiler	71 %
bottleCooker	78 %
average	75%

Table 4.10: Results for 10-fold cross-validation on all task data

Task	Correctly classified
bottteBottleheater	100 %
dishWashing	89%
tea	92%
bottleMicrowave	83%
cupWater	86%
chocolatemilkMicrowave	92%
cleanBottle	81%
soupCooker	95%
bottleWaterboiler	86%
bottleCooker	93%
average	90%

For the multi task data also two groups have been made, bottle tasks and other tasks. The labelling of the data was adjusted, all data was labelled with either, ‘bottle’ or with ‘other’. We trained the classifier on all data (both the single task data and the multi task data). The percentage correctly classified was 88%. The confusion matrix is presented in Table 4.11 and in Table 4.12 the accuracy measures are presented.

Table 4.11: Confusion matrix bottle vs other tasks trained and tested on all task data

bottle	other	<– classified as
71	9	bottle
9	66	other

The true positive rate is for bottle making is 0.89. The false positive rate is 0.12. So in 12% of the cases a task is being classified as bottle making when this was not true.

Conclusion

The results show that the tasks can be recognized in 75 % of the cases. The results are obtained by training on the multi data. When training on the all data, the results are 90%. If

Table 4.12: Accuracy measures bottle vs other tasks

True positive rate	False positive rate	class
0.89	0.12	bottle
0.88	0.11	other

we just focus on the separation of bottle making versus other tasks, then the recognition rate is much higher, 96%. The true positive rate is fairly good (0.89), the false positive rate is high (0.12), and this can be dangerous, because the system will not warn the mother when she has forgotten to prepare a bottle, because the system misclassified other tasks as preparing a bottle.

Chapter 5

Discussion

The research described in this thesis pursued the development of an intelligent monitoring system to support people with a mental disability with their daily life routines. We focussed on a specific problem area, the problem that mothers with a mental disability sometimes forget to give their child a bottle of powdered milk. We want to propose a system that can automatically detect that a bottle has been prepared by the mother, so that the system can give a reminder to the mother when she forgets to prepare a bottle. In this thesis we tried to answer the following research questions:

RQ1. How to monitor that a bottle has been prepared
RQ2. How reliable can it be classified that a bottle for the child has been prepared

The first research question was answered by investigating what sensors to use to monitor that a bottle has been prepared. We analysed which actions are undertaken when preparing a bottle and which objects are used. We equipped these objects with sensors so we could measure when these objects were used. On the basis of this information we could classify that a bottle was being prepared. Measuring object use with different kinds of sensors worked very well. The sensors gave enough information to deduce that an object was used.

The second research question was answered by firstly gathering a data set from experiments. Test subjects were instructed to perform multiple tasks in a test kitchen specially equipped for this research. We gathered two data sets from these experiments. First the test persons were asked to perform the tasks one at the time, and in a second round the subjects were asked to perform multiple tasks at the same time. We obtained two data sets this way, one consisting of single task data, and the other one consisting of the multi task data. We tested three classifiers on these two data sets. A naive Bayesian classifier, a decision tree and a support vector machine. The results for the single task data are very promising, all algorithms scored above the 93 %. When the system was trained on the combination of the single task data and the multi task data it scored 90% correct.
It is important is that the system recognizes that a bottle is being prepared, versus that no bottle has been prepared. We grouped the four tasks of preparing a bottle and made another group of the remaining six tasks. When trained on the single task data the system recognizes in 95% of the cases correctly that a bottle has been prepared and the system never

mis classified other tasks as bottle making.
The same analysis has also been performed on the combination of the single task data and the multi task data. The results showed that in this case in 89% of instances of bottle making were correctly classified, but 12% of the other tasks were wrongfully classified as bottle making. In practice this means that in 12% of the cases the system would not give a warning to the mother when this should happen. This can be dangerous for the child.

To get better results on the multi task data, in future research it can be investigated if another way of representation could work. Our representation did only take into account the total amount of time the objects were used. Other ways of data representation where the sequence of objects use is taken into account could possibly result in higher classification results.
Another possibility is to use a multi agent system where the agents are the classifiers. Confidence levels of the different agents can be used for classifying the tasks.

Things that could be done in the future is extending the system so it can be checked if the milk is prepared in a good way, for example, if the milk is not too hot.
We have already shown that also other tasks such as preparing tea, doing the dishes, warming soup can be detected with this system. The other tasks that can be performed in a kitchen such as making coffee, preparing a meal should also be recognizable. Research to ADLs has already shown that a very broad domain of tasks inside the home can be recognized by using sensors in the objects. Future research could include more tasks, and also unknown tasks for the system. When presenting the system with unknown tasks, confidence levels can be used, so that when a certain confidence level about the correct classification of the task is not reached, the system can indicate that an unknown task is being performed.

In this thesis we have focussed on the first two tasks that an intelligent monitoring system should take over from healthcare workers; monitoring and signaling, which are the input and processing part of the intelligent monitoring system. Future research should concentrate on the output of the system and the two remaining tasks of reporting and intervention. The system should report to the mother that she forgets to feed her child, or maybe the system could also reward her when she prepares a bottle. Further options are to report to the healthcare organisation when the warnings to the mother do not have any effect.
As an intervention it could be thought of that the system would turn on the lights in the kitchen to remind the mother that she has to prepare a bottle.

An important question for future research is how reliable such a system should be. If a system is used in the home of mentally disabled mothers, how many errors is it allowed to make, and what are the consequences when an error occurs. Will the child die from dehydration?
Our opinion is that such a system, that monitors the behaviour of people, should be used as an assistance system, not as a system that controls the people. Therefore the system is not required to be 100 % reliable. The system should function as an extra reminder to the mother, but the full responsibility is with the mother.
Also these systems should never fully take over human intervention.

A healthcare organization could use this research as a basis for a system that can be used in the homes of the mentally disabled people. In the future the system could be tuned for a specific person, this way the system can be more precise, because it can recognize the specific devices used by that person and the specific way of preparing a bottle. The system should be adapted to the special needs of the different mentally disabled clients. This research can be used for other actions than preparing a bottle. For each client is should be decided what actions are important to recognize, in order to assist the client.

The paradigm of recognizing behaviour on the basis of object use can not be used for all kinds of behaviour. It is not possible for example to recognizing if a child is disciplined in a good way. Only behaviour in which specific objects are used can be recognized in this way. For each behaviour or problem to be recognized, an investigation needs to be made about the best way to recognize it.

Information from the sensors can also be used for user profiling. User profiling is keeping a record of the preferences and characteristics of people. By learning the preferences of the different inhabitants of a home, the system can adjust parameters in the home, such as the heating system or sound system. But it can also be used to guide in some routines such as preparing a meal.

Future systems in the home will probably use camera systems to assist the user. Recognizing behaviour on the basis of camera images is developing fastly. One of the domains where the development of recognizing what people are doing is rapidly developing is the game industry. One example is the future system for the Xbox 360, project Natal (Microsoft Corporation, 2009), where no controllers are needed to control the game. The system recognizes the different users and the movements they make on basis of camera images. This could mean that in the future no sensors, except for some cameras, are needed to control our home, for leisure, and to provide safety and assistance to those who need and want it.

Appendix A

General instruction for the subjects

The general instructions for the subjects for the experiment, in Dutch and English.

Instructies proefpersoon

Gedurende het experiment zal je gevraagd worden verschillende opdrachten uit te voeren. Sommige opdrachten hebben betrekking op het klaar maken van een baby flesje. Een stap-voor-stap handleiding zal je hierbij verstrekt worden.

Voor de andere opdrachten geldt dat je die mag uitvoeren zoals je zelf wilt.

Je kunt geen fouten maken.

Als je klaar bent met een opdracht, geef dit dan aan, dan krijg je de volgende opdracht.

In het totaal duurt het experiment ongeveer een uur.

Je krijgt nu de tijd om even rond te kijken in de keuken en te zien waar alles staat.

Laat even weten als er nog vragen zijn.

Alvast bedankt,

Marjolein

Instructions subjects

During the experiment you will be asked to perform different kinds of tasks. Some tasks have to do with the preparation of a bottle for a child. A step-by-step manual will be provided for this.

You can do the other assignments in your own way.

It is not possible to do anything wrong.

Give a sign when you are finished with a task, then a next task will be provided.

The experiment takes about an hour.

You now have the possibility to look around in the kitchen and to see where everything is located.

Please let me know if there are any questions.

Thanks a lot,

Marjolein

Appendix B

Instructions for the subjects about how to prepare a bottle

The instructions for the subjects about how to prepare a bottle, in English and Dutch. There are two versions of the instructions, in the first version the instructions involve first warming the water and then adding the powder milk. In the second version the instructions are to add the powdered milk first and then warm the botte.

Version 1 english

How to prepare a bottle (microwave)

1) half fill the bottle with water
2) put the bottle(without the cap!) in the microwave for 20 seconds on full power
3) add 1 levelled spoon full of milk powder to the water in the bottle
4) put the cap back on the bottle and shake the bottle gently until the milk powder is dissolved in the water

How to prepare a bottle (bottle heater)

1) half fill the bottle with water
2) put the bottle in the bottle heater
3) half fill the bottle heater with water
push the button of the bottle heater to the left
set the kitchen alarm to 2 minutes
6) when the alarm rings, turn off the bottle heater and take the bottle from the bottle heater
7) add one levelled spoon full of milk powder to the water in the bottle
8) put the cap back on the bottle and shake the bottle gently until the milk powder is dissolved in the water

How to prepare a bottle (kettle)

1) fill the kettle with water and turn the kettle on
wait till the kettle has finished or turn the kettle off when the water is warm

wait till the water in the kettle is lukewarm or add some extra cold water
4) half fill the bottle with lukewarm water from the kettle
5) add one levelled spoon full of milk powder to the water in the bottle
6) put the cap back on the bottle and shake the bottle gently until the milk powder is dissolved in the water

How to prepare a bottle (with a pan)

1) fill a pan with water
2) half fill the bottle with water
3) put the bottle upright in the pan with water on the cooker and turn the cooker on
4) set the kitchen alarm to 2 minutes
5) when the alarm rings, turn the cooker off and take the bottle from the pan
6) add one levelled spoon full of milk powder to the water in the bottle
7) put the cap back on the bottle and shake the bottle gently until the milk powder is dissolved in the water

Version 1 dutch

Maken flesvoeding (magnetron)

1) vul het flesje voor de helft met water
2) stop het flesje (zonder dop!) in de magnetron, 20 seconden op vol vermogen
3) voeg 1 afgestreken schepje melkpoeder toe aan het water in de fles
4) draai de dop op de fles en schud licht totdat de melkpoeder is opgelost

Maken flesvoeding (flessenwarmer)

1) vul het flesje tot de helft met water
2) stop het flesje in de flessen warmer
3) vul de flessenwarmer voor de helft met water
4) schuif de knop van de flessenwarmer naar links
5) zet de keuken wekker op 2 minuten
als het alarm afgaat, zet de flessenwarmer uit en haal de fles uit de flessenwarmer
voeg 1 afgestreken schepje melkpoeder toe aan het water in de fles
8) draai de dop op de fles en schud licht totdat de melkpoeder is opgelost

Maken flesvoeding (waterkoker)

1) vul de waterkoker met water en zet de waterkoker aan
2) wacht totdat de waterkoker klaar is of zet hem eerder af als het water warm is
3) wacht totdat het water in de waterkoker handwarm is of doe er extra koud water bij
4) vul het flesje tot de helft met handwarm water uit de waterkoker
5) voeg 1 afgestreken schep melkpoeder toe aan het water in de fles
6) draai de dop op de fles en schud licht totdat de melkpoeder is opgelost

Maken flesvoeding (in de pan)

1) vul een pan met water
2) vul het flesje tot de helft met water
3) zet het flesje rechtop in de pan met water op het kookplaatje en zet het kookplaatje aan
4) zet de keuken wekker op 2 minuten
5) als de wekker gaat, zet het kookplaatje uit en haal het flesje uit het water
6) voeg 1 afgestreken schepje melk toe aan het flesje
7) draai de dop op de fles en schud licht totdat de melkpoeder is opgelost

Version 2 english

How to prepare a bottle (microwave)

1) half fill the bottle with water
2) add 1 levelled spoon full of milk powder to the water in the bottle
3) put the cap back on the bottle and shake the bottle gently until the milk is dissolved in the water
4) put the bottle(without the cap!) in the microwave for 20 seconds on full power
5) take the bottle from the microwave and put the cap back on the bottle. Shake the bottle gently so the warmth is being spread equally

How to prepare a bottle (bottle heater)

1) half fill the bottle with water
2) add one levelled spoon full of milk powder to the water in the bottle
3) put the cap back on the bottle and shake the bottle gently until the milk powder is dissolved in the water
4) put the bottle in the bottle heater
5) half fill the bottle heater with water
6) push the button of the bottle heater to the left
7) set the kitchen alarm to 2 minutes
8) when the alarm rings, turn off the bottle heater and take the bottle from the bottle heater
9) shake the bottle gently so the warmth is being spread equally

How to prepare a bottle (kettle)

1) fill the kettle with water and turn the kettle on
2) wait till the kettle has finished or turn the kettle off when the water is warm
3) wait till the water in the kettle is lukewarm or add some extra cold water
4) half fill the bottle with lukewarm water from the kettle
5) add one levelled spoon full of milk powder to the water in the bottle
6) put the cap back on the bottle and shake the bottle gently until the milk powder is dissolved in the water

How to prepare a bottle (with a pan)

1) fill a pan with water
2) half fill the bottle with water
3) add one levelled spoon full of milk powder to the water in the bottle
4) put the cap back on the bottle and shake the bottle gently until the milk is dissolved in the water
5) put the bottle upright in the pan with water on the cooker and turn the cooker on
6) set the kitchen alarm to 2 minutes
7) when the alarm rings, turn the cooker off and take the bottle from the pan

8) shake the bottle gently so the warmth is being spread equally

Version 2 dutch

Maken flesvoeding (magnetron)

1) vul het flesje voor de helft met water
2) voeg 1 afgestreken schepje melkpoeder toe
3) doe de dop op het flesje en schud het flesje licht totdat de melkpoeder is opgelost in het water
4) stop het flesje (zonder dop!) in de magnetron, 20 seconden op vol vermogen
5) haal het flesje uit de magnetron en doe de dop weer op het flesje. Schud het flesje licht totdat de warmte goed verdeelt is

Maken flesvoeding (flessenwarmer)

1) vul het flesje voor de helft met water
2) voeg 1 afgestreken schepje melkpoeder toe
3) doe de dop op het flesje en schud het flesje licht totdat de melkpoeder is opgelost in het water
4) stop het flesje in de flessenwarmer
5) vul de flessenwarmer voor de helft met water
6) schuif de knop van de flessenwarmer naar links
7) zet de keuken wekker op 2 minuten
8) als het alarm afgaat, zet de flessenwarmer uit en haal de fles uit de flessenwarmer
9) schud het flesje licht totdat de warmte goed verdeelt is

Maken flesvoeding (waterkoker)

1) vul de waterkoker met water en zet de waterkoker aan
2) wacht totdat de waterkoker klaar is of zet hem eerder af als het water warm is
3) wacht totdat het water in de waterkoker handwarm is of doe er extra koud water bij
4) vul het flesje tot de helft met handwarm water uit de waterkoker
5) voeg 1 afgestreken schep melkpoeder toe aan het water in de fles
6) draai de dop op de fles en schud licht totdat de melkpoeder is opgelost

Maken flesvoeding (in de pan)

1) vul een pan met water
2) vul een flesje tot de helft met water
3) voeg 1 afgestreken schepje melk toe
4) draai de dop op de fles en schud licht totdat de melkpoeder is opgelost
5) zet de fles rechtop in de pan met water op het kookplaatje en zet het kookplaatje aan
6) zet de keuken wekker op 2 minuten
7) als de wekker gaat, zet het kookplaatje uit en haal het flesje uit het water
8) schud het flesje licht totdat de warmte goed verdeelt is

Appendix C

Task cards

Prepare some tea	Warm a cup of hot chocolate in the microwave
Prepare a bottle using the microwave	Prepare some soup on the cooker
Prepare a bottle using the kettle	Clean the bottle and put it back in the kitchen cabinet
Prepare a bottle using a pan	clean the cups and place them back in the kitchen cabinet
Prepare a bottle using the bottle heater	grab a cup of water

Avery C32011

Figure C.1: Instruction cards english

Appendix D

Result data

D.1 Single task classification

D.1.1 10-fold cross-validation

Table D.1: Results single task classification for 10-fold cross-validation

Algorithms	percentage correctly classified
SMO	93.6%
C4.5	92.7%
Naive Bayes	93.6%

Table D.2: Confusion matrix SMO algorithm single task classification 10-fold cross-validation

a	b	c	d	e	f	g	h	i	j	<– classified as
11	0	0	0	0	0	0	0	0	0	a = bottleBottleheater
0	9	0	0	1	0	1	0	0	0	b = dishWashing
0	0	11	0	0	0	0	0	0	0	c = tea
0	0	0	10	0	0	1	0	0	0	d = bottleMicrowave
0	0	0	0	11	0	0	0	0	0	e = cupWater
0	1	0	0	0	10	0	0	0	0	f = chocolademilkMicrowave
0	2	0	0	0	0	9	0	0	0	g = cleanBottle
0	0	0	0	0	0	0	11	0	0	h = soupCooker
0	0	1	0	0	0	0	0	10	0	i = bottleWaterboiler
0	0	0	0	0	0	0	0	0	11	j = bottleCooker

Table D.3: Confusion matrix C4.5 algorithm single task classification 10-fold cross-validation

a	b	c	d	e	f	g	h	i	j	<– classified as
11	0	0	0	0	0	0	0	0	0	a = bottleBottleheater
0	10	0	0	1	0	0	0	0	0	b = dishWashing
0	0	10	0	0	0	0	1	0	0	c = tea
0	0	0	10	0	1	0	0	0	0	d = bottleMicrowave
0	2	0	0	9	0	0	0	0	0	e = cupWater
0	1	0	0	0	9	1	0	0	0	f = chocolademilkMicrowave
0	0	0	0	0	0	10	0	1	0	g = cleanBottle
0	0	0	0	0	0	0	11	0	0	h = soupCooker
0	0	0	0	0	0	0	0	11	0	i = bottleWaterboiler
0	0	0	0	0	0	0	0	0	11	j = bottleCooker

Table D.4: Confusion matrix Naive Bayes algorithm single task classification 10-fold cross-validation

a	b	c	d	e	f	g	h	i	j	<– classified as
11	0	0	0	0	0	0	0	0	0	a = bottleBottleheater
0	9	0	0	1	0	1	0	0	0	b = dishWashing
0	0	10	0	0	0	0	0	1	0	c = tea
0	0	0	11	0	0	0	0	0	0	d = bottleMicrowave
0	0	0	0	11	0	0	0	0	0	e = cupWater
0	0	0	1	0	10	0	0	0	0	f = chocolademilkMicrowave
0	0	0	0	0	0	11	0	0	0	g = cleanBottle
0	0	0	0	0	0	0	9	0	2	h = soupCooker
0	0	0	1	0	0	0	0	10	0	i = bottleWaterboiler
0	0	0	0	0	0	0	0	0	11	j = bottleCooker

D.1.2 leave-one-subject-out cross-validation

Table D.5: Results single task classification for 10-fold cross-validation

Algorithms	percentage correctly classified
SMO	94.5%
C4.5	93.6%
Naive Bayes	92.7%

Table D.6: Results single task classification for the SMO algorithm with leave-one-subject-out cross-validation

Test person	percentage correctly classified
1	90%
2	100%
3	80%
4	90%
5	100%
6	100%
7	100%
8	100%
9	90%
10	100%
11	90%

Table D.7: Confusion matrix SMO algorithm single task classification leave-one-subject-out cross-validation for subject 1

a	b	c	d	e	f	g	h	i	j	<– classified as
1	0	0	0	0	0	0	0	0	0	a = bottleBottleheater
0	0	0	0	1	0	0	0	0	0	b = dishWashing
0	0	1	0	0	0	0	0	0	0	c = tea
0	0	0	1	0	0	0	0	0	0	d = bottleMicrowave
0	0	0	0	1	0	0	0	0	0	e = cupWater
0	0	0	0	0	1	0	0	0	0	f = chocolademilkMicrowave
0	0	0	0	0	0	1	0	0	0	g = cleanBottle
0	0	0	0	0	0	0	1	0	0	h = soupCooker
0	0	0	0	0	0	0	0	1	0	i = bottleWaterboiler
0	0	0	0	0	0	0	0	0	1	j = bottleCooker

Table D.8: Confusion matrix SMO algorithm single task classification leave-one-subject-out cross-validation for subject 3

a	b	c	d	e	f	g	h	i	j	<– classified as
1	0	0	0	0	0	0	0	0	0	a = bottleBottleheater
0	1	0	0	0	0	0	0	0	0	b = dishWashing
0	0	1	0	0	0	0	0	0	0	c = tea
0	0	0	1	0	0	0	0	0	0	d = bottleMicrowave
0	0	0	0	1	0	0	0	0	0	e = cupWater
0	1	0	0	0	0	0	0	0	0	f = chocolademilkMicrowave
0	1	0	0	0	0	0	0	0	0	g = cleanBottle
0	0	0	0	0	0	0	1	0	0	h = soupCooker
0	0	0	0	0	0	0	0	1	0	i = bottleWaterboiler
0	0	0	0	0	0	0	0	0	1	j = bottleCooker

Table D.9: Confusion matrix SMO algorithm single task classification leave-one-subject-out cross-validation for subject 4

a	b	c	d	e	f	g	h	i	j	<– classified as
1	0	0	0	0	0	0	0	0	0	a = bottleBottleheater
0	1	0	0	0	0	0	0	0	0	b = dishWashing
0	0	1	0	0	0	0	0	0	0	c = tea
0	0	0	1	0	0	0	0	0	0	d = bottleMicrowave
0	0	0	0	1	0	0	0	0	0	e = cupWater
0	0	0	0	0	1	0	0	0	0	f = chocolademilkMicrowave
0	0	0	0	0	0	1	0	0	0	g = cleanBottle
0	0	0	0	0	0	0	1	0	0	h = soupCooker
0	0	1	0	0	0	0	0	0	0	i = bottleWaterboiler
0	0	0	0	0	0	0	0	0	1	j = bottleCooker

Table D.10: Confusion matrix SMO algorithm single task classification leave-one-subject-out cross-validation for subject 9

a	b	c	d	e	f	g	h	i	j	<– classified as
1	0	0	0	0	0	0	0	0	0	a = bottleBottleheater
0	1	0	0	0	0	0	0	0	0	b = dishWashing
0	0	1	0	0	0	0	0	0	0	c = tea
0	0	0	1	0	0	0	0	0	0	d = bottleMicrowave
0	0	0	0	1	0	0	0	0	0	e = cupWater
0	0	0	0	0	1	0	0	0	0	f = chocolademilkMicrowave
0	1	0	0	0	0	0	0	0	0	g = cleanBottle
0	0	0	0	0	0	0	1	0	0	h = soupCooker
0	0	0	0	0	0	0	0	1	0	i = bottleWaterboiler
0	0	0	0	0	0	0	0	0	1	j = bottleCooker

D.2 Multi task classification

Maak thee	Warm een beker chocolade melk op in de magnetron
Maak een flesje met behulp van de magnetron	Maak soep klaar op het kookplaatje
Maak een flesje met behulp van de waterkoker	Maak het flesje schoon en zet het terug in het keukenkastje
Maak een flesje met behulp van een pannetje water	Maak de kopjes schoon en zet ze terug in het keukenkastje
Maak een flesje klaar met behulp van flessenwarmer	pak een beker water

Avery C32011

Figure C.2: Instruction cards dutch

Table D.11: Confusion matrix SMO algorithm single task classification leave-one-subject-out cross-validation for subject 11

a	b	c	d	e	f	g	h	i	j	<– classified as
1	0	0	0	0	0	0	0	0	0	a = bottleBottleheater
0	1	0	0	0	0	0	0	0	0	b = dishWashing
0	0	1	0	0	0	0	0	0	0	c = tea
0	0	0	0	0	0	1	0	0	0	d = bottleMicrowave
0	0	0	0	1	0	0	0	0	0	e = cupWater
0	0	0	0	0	1	0	0	0	0	f = chocolademilkMicrowave
0	0	0	0	0	0	1	0	0	0	g = cleanBottle
0	0	0	0	0	0	0	1	0	0	h = soupCooker
0	0	0	0	0	0	0	0	1	0	i = bottleWaterboiler
0	0	0	0	0	0	0	0	0	1	j = bottleCooker

Table D.12: Results for 10-fold cross-validation on multi task data

Task	Correctly classified
bottteBottleheater	100%
dishWashing	76 %
tea	69 %
bottleMicrowave	67 %
cupWater	73 %
chocolatemilkMicrowave	80 %
cleanBottle	60 %
soupCooker	78 %
bottleWaterboiler	71 %
bottleCooker	78 %
average	75%

Table D.13: Results for 10-fold cross-validation on all task data

Task	Correctly classified
bottteBottleheater	100 %
dishWashing	89%
tea	92%
bottleMicrowave	83%
cupWater	86%
chocolatemilkMicrowave	92%
cleanBottle	81%
soupCooker	95%
bottleWaterboiler	86%
bottleCooker	93%
average	90%

Table D.14: Confusion matrix bottle vs other tasks trained and tested on all task data

bottle	other	<– classified as
71	9	bottle
9	66	other

Table D.15: Accuracy measures bottle vs other tasks

True positive rate	False positive rate	class
0.88	0.12	bottle
0.88	0.11	other

Bibliography

Akyildiz, I. F., Su, W., Sankarasubramaniam, Y., & Cayirci, E. (2002). Wireless sensor networks: a survey. *Computer Networks*, *38*(4), 393-422.

Alwan, M., Leachtenauer, J., Dalal, S., Kell, S., Turner, B., Mack, D., et al. (2005). Validation of rule-based inference of selected independent activities of daily living. *Telemedicine and e-Health*, *11*(5), 594-599.

American Psychiatric Association. (1994). *Diagnostic and statistical manual of mental disorders* (4 ed.). Washington, DC, USA: American Psychiatric Association.

Argandoña, E. S. de, Aztiria, A., García, C., Arana, N., Izaguirre, A., & Fillatreau, P. (2008). Forming processes control by means of artificial intelligence techniques. *Robotics and Computer-Integrated Manufacturing*, *24*(6), 773-779.

Baker, C. R., Armijo, K., Belka, S., Benhabib, M., Bhargava, V., Burkhart, N., et al. (2007). Wireless sensor networks for home health care. In *AINAW '07: Proceedings of the 21st International Conference on Advanced Information Networking and Applications Workshops* (pp. 832–837). Washington, DC, USA: IEEE Computer Society.

Bakker, R. (2008). *Ondersteuning van verstandelijk gehandicapten door slimme (sensor)technologie.* (Project proposal, personal communications, june 11, 2009)

Bao, L., & Intille, S. S. (2004). Activity recognition from user-annotated acceleration data. In A. Ferscha & F. Mattern (Eds.), *Second international conference on pervasive computing* (Vol. 3001, p. 1-17). Berlin / Heidelberg: Springer.

Burrell, J., Brooke, T., & Beckwith, R. (2004). Vineyard computing: Sensor networks in agricultural production. *IEEE Pervasive Computing*, *3*(1), 38–45.

Butler, Z., Corke, P., Peterson, R., & Rus, D. (2004). Dynamic virtual fences for controlling cows. In *Experimental robotics IX* (Vol. 21/2006, p. 513-522). Berlin / Heidelberg: Springer.

Chan, M., Esteve, D., Escriba, C., & Campo, E. (2008). A review of smart homes—present state and future challenges. *Computer methods and programs in biomedicine*, *91*, 55-81.

Cook, D. J., Youngblood, M., & Das, S. K. (2006). A multi-agent approach to controlling a smart environment. In J. C. Augusto & C. D. Nugent (Eds.), *Designing smart homes* (Vol. 4008/2006, p. 165-182). Berlin / Heidelberg: Springer.

Dalal, S., Alwan, M., Seifrafi, R., Kell, S., & Brown, D. (2005). *A rule-based approach to the analysis of elders activity data: Detection of health and possible emergency conditions.* Retrieved November 5, 2008, from https://www.aaai.org/Papers/Symposia/Fall/2005/FS-05-02/FS05-02-006.pdf.

Damas, M., Prados, A., Gómez, F., & Olivares, G. (2001). Hidrobus® system: fieldbus for integrated management of extensive areas of irrigated land. *Microprocessors and Microsystems*, *25*, 177-184.

Gao, T., Greenspan, D., Welsh, M., Juang, R., & Alm, A. (2005). Vital signs monitoring and patient tracking over a wireless network. In *Proceedings of the 27th annual international conference of the ieee emes.* Shanghai: Johns Hopkins University Applied Physics Laboratory, Laurel, MD.

Harsányi, G., Bojta, P., Gordon, P., Lepsényi, I., & Ballun, G. (2009). *SensEdu.* Retrieved July 8, 2009, from http://www.sensedu.com.

Huang, K.-T., Lin, P.-Y., Chiang, C.-Y., Chang, J.-S., Huang, C.-N., & Chan, C.-T. (2008). An intelligent RFID system for improving elderly daily life independent in indoor environment. In *Smart homes and health telematics* (p. 1-8). Berlin / Heidelberg: Springer.

Huỳnh, T., Blanke, U., & Schiele, B. (2007). Scalable recognition of daily activities with wearable sensors. In *Location- and context-awareness* (Vol. 4718/2007, p. 50-67). Berlin / Heidelberg: Springer.

INCAS^3. (2009). *INCAS3: Sensor Concept.* Retrieved July 30, 2009, from http://www.incas3.eu/lang/en/research/sensor-concept.

Isoda, Y., Kurakake, S., & Nakano, H. (2004). Ubiquitous sensors based human behavior modeling and recognition using a spatio-temporal representation of user states. In *AINA '04: Proceedings of the 18th international conference on advanced information networking and applications.* Washington, DC, USA: IEEE Computer Society.

Keerthi, S., Shevade, S., Bhattacharyya, C., & Murthy, K. (2001). Improvements to Platt's SMO Algorithm for SVM Classifier Design. *Neural Computation*, *13*(3), 637-649.

Logan, B., Healey, J., Philipose, M., Tapia, E. M., & Intille, S. (2007). A long-term evaluation of sensing modalities for activity recognition. In *UbiComp 2007 : Ubiquitous Computing* (Vol. 4717/2007, p. 483-500). Berlin / Heidelberg: Springer.

Lorincz, K., Malan, D. J., Fulford-Jones, T., Nawoj, A., Clavel, A., Shnayder, V., et al. (2004). Sensor networks for emergency response: Challenges and opportunities. *Pervasive computing*, *3*(4), 16-23.

Meulen, J. van der, Schaaf, M. van der, Tan, J., & Wegman, C. (2008). *Project: "jonge moeder".* (Project report CENSI, personal commmunications, september 9, 2008)

Microsoft Corporation. (2009). *Project NATAL.* Retrieved August 4, 2009, from http://www.xbox.com/en-US/live/projectnatal.

Milenkovic, A., Otto, C., & Jovanov, E. (2006). Wireless sensor networks for personal health monitoring: Issues and an implementation. *Wirelsess Senson Networks and Wired/Wireless Internet Communications*, *29*(13-14), 2521-2533.

Nagl, L., Schmitz, R., Warren, S., Hildreth, T., Erickson, H., & Andresen, D. (2003). *Wearable sensor system for wireless state-of-health determination in cattle.* Retrieved July 8, 2009, from http://people.cis.ksu.edu/ dan/Vet_Cancun_BMOO_Final.pdf.

Patterson, D. J., Fox, D., Kautz, H., & Philipose, M. (2005). Fine-grained activity recognition by aggregating abstract object usage. In *ISWC '05: Proceedings of the Ninth IEEE International Symposium on Wearable Computers* (p. 44-51). Washington, DC, USA: IEEE Computer Society.

Phidgets, Inc. (2008). *Phidgets Inc. - unique and easy to use USB interfaces.* Retrieved March 13, 2009, from http://www.phidgets.com.

Philipose, M., Fishkin, K. P., Perkowitz, M., Patterson, D. J., Fox, D., Kautz, H., et al. (2004). Inferring activities from interactions with objects. *Pervasive computing*, *3*(4), 50-57.

Platt, J. (1998). Machines using sequential minimal optimization. In B. Schoelkopf, C. Burges, & A. Smola (Eds.), *Advances in kernel methods - support vector learning* (p. 41-65). Cambridge, MA: MIT Press.

Quinlan, R. (2003). *C4.5: Programs for Machine Learning.* San Matea, CA: Morgan Kaufmann Publishers.

Russell, S. J., & Norvig, P. (2003). *Artificial intelligence: a modern approach* (2nd ed.). Upper Saddle River, NJ: Prentice Hall/Pearson Education.

Sarle, W. (2009). *Should I normalize/standardize/rescale the data.* Retrieved July 26, 2009, from http://www.faqs.org/faqs/ai-faq/neural-nets/part2/section-16.html.

Soucek, S., Russ, G., & Tamarit, C. (2000, December). *The smart kitchen project – an application of fieldbus technology to domotics.* Retrieved July 9, 2009, from http://smartkitchen.ict.tuwien.ac.at/publications/iwna2000.pdf. New Brunswick, NJ, USA.

Tapia, E. M., Intille, S. S., & Larson, K. (2004). *Activity recognition in the home using simple and ubiquitous sensors.* Retrieved September 3, 2008, from http://www.springerlink.com/content/5a4qm20y37089gk9. Berlin Heidelberg: Springer-Verlag.

Thebe. (2009). *GezondVGZ - Flesvoeding klaarmaken.* Retrieved July 14, 2009, from http://www.gezondvgz.nl/d_cm7004.

Uraikul, V., Chan, C. W., & Tontiwachwuthikul, P. (2007, March). Artificial intelligence for monitoring and supervisory control of process systems. *Engineer Applications of Artificial Intelligence*, *20*(2), 115-131.

Velera, M., & Velastin, S. (2005, April). Intelligent distributed surveillance systems: a review. *IEE Proceedings - Vision, Image and Signal Processing*, *152*(2).

Vivoni, E. R., & Camilli, R. (2003, May). Real-time streaming of environmental field data. *Computers and Geosciences*, *29*(4), 457-468.

Weka. (2009). *Weka 3 - Data Mining with Open Source Machine Learning Software in Java.* Retrieved May 20, 2009, from http://www.cs.waikato.ac.nz/ml/weka.

Wilson, D., & Atkeson, C. (2005). Simultaneous tracking & activity recognition (STAR) using many anonymous, binary sensors. In *Pervasive computing* (p. 62-79). Berlin / Heidelberg: Springer.

Witten, I. H., & Frank, E. (2005). *Data mining: practical machine learning tools and techniques* (2nd ed.). San Francisco, CA: Morgan Kaufmann Publishers.

Zhang, H. (2004). The optimality of naive Bayes. In *Proceedings of the seventeenth international florida artificial intelligence researc society conference.* Miami Beach, Florida, USA: AAAI Press.