Exploring The Responsibilities Of Single-Inhabitant Smart Homes With Use Cases

Paul Lyons[1], An Tran Cong, H. Joe Steinhauer, Stephen Marsland, Jens Dietrich, and Hans W. Guesgen
School of Engineering and Advanced Technology
Massey University
Palmerston North
New Zealand

[1] *Ph: +6463569099 x 2472*
Fax: +6463502259
Email: P.Lyons@massey.ac.nz

Abstract. This paper makes a number of contributions to the field of requirements analysis for Smart Homes. It introduces Use Cases as a tool for exploring the responsibilities of Smart Homes and it proposes a modification of the conventional Use Case structure to suit the particular requirements of Smart Homes. It presents a taxonomy of Smart-Home-related Use Cases with seven categories. It draws on those Use Cases as raw material for developing questions and conclusions about the design of Smart Homes for single elderly inhabitants, and it introduces the SHMUC repository, a web-based repository of Use Cases related to Smart Homes that anyone can exploit and to which anyone may contribute.

Keywords: Use Case, taxonomy, Smart Home, Smart Environments, SHMUC.

1. Introduction

The aim of the Massey University Smart Environment (MUSE) project [4, 5] is to develop unobtrusive, affordable, computationally inexpensive, off-the-shelf Smart Home technologies for use in elder-care, particularly Smart Homes that cater for a single inhabitant. It is important that the designed system fulfils the requirements that the elderly person living in it, their relatives relying on it, and the carers depending on it, have for such a system. The ideas that are documented here derive from discussions within the MUSE group about Smart Homes with a single elderly inhabitant. They may have application in other areas, but they are primarily related to that particular universe of discourse.

One motivation for designing such a system is to enable the elderly to live in a familiar place as long as possible. In particular elderly people who suffer from cognitive impairment are known to achieve a higher quality of life and remain independent longer when living in their own home. The reason is that many tasks fulfilled on a daily basis, often called ADLs (Activities of Daily Living), such as eating, dressing and grooming [11] are over-learned and automated processes that the elderly can still perform if they remain in a place where they are used to performing them, but are fragile when they have to be performed in unfamiliar places [1].

As the population ages [17], the importance of developing caring environments for the elderly will inevitably grow and there are a number of related projects, such as Adaptive House [18], iDorm [6], MavHome [21], Georgia Tech Aware Home [9], PlaceLab [16], and Gator Tech Smart House [7] dealing with various aspects of Smart Homes for the elderly.

A good deal of the challenge in developing such systems resides in their interaction with humans. It is not enough simply to provide useful functions; the functions have to integrate well into the daily lives of the people who inhabit the Smart Homes. One technique for teasing out the requirements for such complex applications is to construct a number of Use Cases [2, 8], each of which presents a single scenario – a realistic example of a situation in which the application will be used. This technique has been extensively used in software engineering because it helps designers to focus their effort on functions that real users will need and ways in which real users will interact with the system. It often reveals functions whose need might otherwise not have become apparent until later, when the application was in production.

We have adopted this approach for analyzing the requirements of Smart Homes, and have adapted the conventional form of Use Cases to the particular, and somewhat unusual, requirements of this application area. This paper describes and explains those adaptations. It also presents a number of Smart Home-related Use Cases, and outlines some conclusions and questions that the exercise has thrown up. A tentative preliminary taxonomy of Use Cases with four major categories and seven leaf nodes is also illustrated.

In addition, the paper introduces SHMUC, a freely editable web-based repository of Use Cases which is intended to stimulate and record discussions by the wider smart environment community of requirements for Smart Homes, and to act as a requirements resource for Smart Home developers. Although their application area is quite different, The World Wide Web consortium have also exploited the idea of a web-based repository of Use Cases to assist it gathering and integrating the opinions of interested parties when developing W3C recommendations [19].

It should be noted that the philosophy of the SHMUC repository is that the Use Cases it contains are always under development and neither the repository nor the Use Cases presented in this paper are the last word on the subject; although inconsistencies are deplorable, they are not fatal. There are several types of technology that might be incorporated into a Smart Home; The Gator Tech Smart House [7] is an example of a high-tech system that incorporates intelligent appliances such as smart beds, smart floors, and smart washing machines. However, such high-tech approaches have social implications. They are expensive, and it is not only the rich who deserve the dignity of an independent lifestyle in their old age. We are therefore following the example of PlaceLab [16], and opting for a more inexpensive solution where simple sensors, such as state change sensors and accelerometers which can be attached to as many appliances as possible in the house. The advantages are that the sensors are relatively inexpensive, they can be quickly installed, and the activities are only indirectly observed so the inhabitant's privacy is assured.

The overall motivation for the Use Cases presented here is to develop a tool for identifying abnormal behaviors and suitable ways of dealing with them. Before considering that tool, however, it is worth considering the nature of abnormality and the particular problems that it engenders.

The universal set of human behaviors is intractably large, as is the subset of interest, the set of abnormal behaviors. If a complete set of normal behaviors were available to the Smart Home, it could identify abnormal behaviors by a trivial application of set theory; any behaviors that were not members of the set of normal behaviors would *ipso facto* be abnormal. Unfortunately, the complete set of normal behaviors, although smaller, is also so large as to be infeasible to collect.

However, it is conceivable that, with a considerable amount of effort, a significant number of the normal behaviors of a population of subjects could be documented. A Smart Home that was seeded with this subset of normal behaviors would be able to identify, and allow to continue without let or hindrance, a significant proportion of the inhabitant's normal behaviors. An abnormality detector that worked by complementing this subset of normal behaviors would inevitably produce annoying false positives, and a user-friendly Smart Home would need to be able to reduce their frequency by adding to its set of normal behaviors. Ideally, it would do this by learning about new normal behaviors for itself, but more likely by having its model corrected by a human when errors occurred. This SUbset of Normal Behaviors, Augmenting Model (SUNBEAM) is the model of abnormality detection that underlies the work presented here.

However, abnormality is not a sufficient criterion for intervention by the Smart Home. We have introduced the concepts of *interesting* and *problematic* behaviours. All abnormal behaviors are considered to be *interesting*, and worthy of further examination. Only if the result of this examination is that the behavior is also classified as *problematic* does the system raise an alert. Finally, the appropriate form of intervention is decided upon.

Use Cases were not originally developed for use in abnormality detection, and, as we shall see, a modified version of the format has been developed to suit them to this application.

2. Approach

It is not easy to pin down the responsibilities of a Smart Home. The field is broad, and the range of activities that is encompassed by the term Smart Home is correspondingly extensive. Many, many realistic scenarios have emerged during the MUSE research group's discussions in an effort to ensure that the – sometimes quite abstruse – techniques that were being proposed for analysis of human behavior were grounded in reality. However, although these informal scenarios often allowed us to examine facets of Smart Home design through a high-powered microscope, they did not make it easy to stand back and obtain a wider field of view. Indeed, some meetings degenerated into war-by-scenario; one group member would propose a scenario that supported one analytic approach, and another group member would immediately propose a minimally different scenario that supported an alternative approach. What was required was a technique that, like our informal scenarios, was grounded in approachable, realistic examples and, at the same time, made it easy to distinguish and extricate general principles from an excess of detail. At some point, it was suggested that presenting the examples as formal Use Cases would help us satisfy this need. Each Use Case would codify the system's behavior in response to a specific user goal. It was hoped that, if a sufficient population of Use Cases was built up, it would be possible to stand back, survey them all, group similar Use Cases, and extract the common elements to form a comparatively small number of general categories of behavior. Each of these categories of behavior could subsequently be incorporated as a function in the Smart Home's capability set, without the need to write code to handle each situation explicitly.

The range of activities encompassed by the term Smart Home is extensive, even when the focus is on highly specialized homes that can give a person, generally an older person, whose physical or mental faculties have diminished, unobtrusive support that allows them to maintain an independent lifestyle for longer than would otherwise be possible. Such a system has some particular characteristics that make it a particularly challenging design task:

- The Smart Home needs to be able to analyze and make inferences about human behavior.
- The inhabitant of the Smart Home is an unusual "user;" she or he is trying to live as independent a life as possible, and may prefer never to interact directly with the Smart Home at all.
- Although the inhabitant may be pursuing explicit goals, and although the pursuit of those goals may cause or require the system to react, the goals that drive the system's behavior may be distinct from, or even in contradiction to, the inhabitant's goals.

An example of the last characteristic might occur when the inhabitant starts to cook a meal, forgets about it, and doesn't turn the stove off. This could lead to the food, and then the house, catching fire. A thoroughly competent Smart Home would detect that the temperature of the air above the stove had risen to 200C, and (to anthropomorphize only a little), intervene to turn the stove off. In this case, the inhabitant's (forgotten) goal was to cook a meal, whereas the Smart Home's goal was to prevent the house from burning down.

We have started to use Uses Cases as an exploratory tool to help uncover the implications of the phrase "a thoroughly competent Smart Home." Such an approach is particularly apposite as, at the time of writing, the research group is attempting to decide what a Smart Home for a single elderly inhabitant should be responsible for; the group's activities are therefore largely focused on requirements analysis. Use Cases are widely used in software engineering at the requirements analysis stage of designing systems [14]. They are built around readily comprehensible real-life scenarios that make it easy to distil the essence of the system and whose significance will be readily understood without recourse to technical jargon. We believe that Use Cases may be a suitable tool for exploring the capabilities of a more general class of smart environment than the single-elderly-inhabitant Smart Homes that are the particular focus of the MUSE research group, and that smart environment researchers in general could benefit from their adoption. Although the examples presented here, and the conclusions they support, are specific to that area of investigation, we hope that it will be clear that other Use Cases could be developed and used to illustrate the requirements of other types of smart environment.

Use Cases are attractive for two reasons. On one hand they are readily understood by stakeholders who are not experts in software development or engineering in general, because they are grounded in real human experience (although the "stories" they contain are usually fabricated). On the other hand, they are a blueprint for system designers, because each Use Case focuses on one aspect of a system's behavior, and is a clear exposition of the responsibilities of a code module in the resulting system.

What's in a Use Case? A Use Case is a form of structured English; although a number of diagrammatic representations have been proposed, it generally comprises a number of text fields of greater or lesser extent. The structure is "standardised," but in fact many structures for Use Cases have been promulgated. Cockburn [2] cites 18. We have followed the trend and developed our own form, which is true to the Jacobson's original intent, but has a number of adaptations that fit it to the specialized nature of the Smart Home domain. However, before we describe those modifications, let us consider the general structure of a Use Case.

The foundation on which the rest of the structure sits is a description of an interaction between a system and a number of "actors." The actors are most often humans, who are given names to make them seem realistic to readers (and perhaps equally importantly, to subtly influence the author of the Use Case towards writing actions that real people might perform), but institutions may be involved, and so may other computer systems. In the end, the category is so broad that it even makes sense to portray the system itself as one of the actors. The interaction at the centre of the Use Case generally involves an unbranched, partially ordered sequence of actions, usually alternating between a principal actor and the system. The Use Case documents the principal actor's goal in interacting with the system, and includes a short story, written in jargon-free language, that outlines the interaction from the principal actor's point of view.

Here is a conventionally structured Use Case for programming home heating over a cell phone:

Use Case:
Heat the house before arriving home from holiday.

Actors:
Billy, and the Smart Home

Goal:
Billy wishes to come home to a warm house after a winter holiday in the tropics

Scenario:
Billy has taken a month off work during the winter to go on holiday to the Caribbean. The holiday hasn't turned out well – his girlfriend dumped him on the last day – and on disembarking at his small town airport from the 30-seater plane that has barely made it back to land, he finds himself buffeted by howling wind and freezing rain that is "falling" horizontally across the runway. Billy is miserable and he still faces an hour's bus ride to get home. However, he remembers that there is one thing in his favor. He can turn on the hot water heating and the central heating at home, so that when he finally gets home, his house will be warm and he will be able to have a long hot shower.

Interaction Sequence:

1	Billy:	phones home
2	Home:	asks the caller for ID
3	Billy:	identifies himself[1]

[1] Note that this is deliberately vague. In the current technological environment, Billy probably uses a PIN to identify himself, but it is possible that retinal scan, fingerprint identification, or some other mechanism

4	Home:	prompts Billy to choose an action
5	Billy:	chooses *Water heating*
6	Home:	gives Billy a choice between *On: eco-sensitive*, *On: eco-destructive* and *Off* (third option not selectable)
7	Billy:	chooses *On: eco-destructive*
8	Home:	turns on the water heater prompts Billy to choose an action
9	Billy:	chooses *Central heating*
10	Home:	gives Billy a choice between *On: low, On: medium* and *On: hot, hot, hot*
11	Billy:	chooses *On: hot, hot, hot*
12	Home:	prompts Billy to choose an action
13	Billy:	chooses *Log out*
		etc.

This Use Case makes it clear what type of options are available to Billy, and also makes it clear to the designer of the software that his Smart Home doesn't just give him a single option and then hang up, forcing him to call up repeatedly if he wishes to control more than one appliance. This second requirement for the system could be one that emerged as a result of writing this Use Case; it was not foreseen when the Use Case was first outlined.

Now, this is not a complete Use Case. It only contains the Success Scenario [2], which describes what happens if everything goes well. But, of course, the scenario might not complete successfully. For example, Billy might fail to identify himself correctly, or the central heating might fail to start up. A Use Case may incorporate one or more Exception Scenarios, which are invoked if a step in the Success Scenario cannot be completed. Logically, the structure of the Use Case is a tree, with a branch at each point where the system may depart from the sequential list of actions in the Success Scenario. Typographically, it is represented as a set of lists; the Success Scenario is one list, and each associated Exception Scenario is written as a separate list of actions. This allows for non-specialist stakeholders who may not be capable of dealing with the complexity of a tree traversal.

As we shall see, the exception scenario concept does not map well onto the Smart Homes designed for elder-care, and most of the Use Cases we list do not include Exception Scenarios.

The Use Case approach treads a fine line between being informal enough to be capable of communicating ideas to, and extracting opinions from, non-specialists, and being formal enough to act as the starting-point for a system specification. For example, the steps in an interaction sequence may be viewed as subgoals with their own Use Cases, so that in general, Use Cases are recursive structures. This mathematically elegant but, for many people, intellectually challenging structure does not have to be pointed out to clients, but pre-prepared sub-Use Cases can be brought into play should the clients query the completeness of the high-level representation.

However, there are two areas of mismatch between conventional uses cases, as shown in the central heating example above and the Smart Home application area; goals and the dialogs that occur in interaction sequences.

A conventional description of the purpose of Use Cases might say that they document the information transfers that occur when a user initiates an interaction that is intended to achieve a particular goal. Two aspects of this description are true in general, but may not apply to interactions with a Smart Home. First, it is implicit in the description that interactions between a user and a system occur at the behest of a user, and second that, when an interaction between a user and a system occurs, the user has a particular goal in mind. For the Smart Homes that concern the MUSE research group, both of these conditions may be false. In the context of our research, the primary "users" of Smart Homes are their elderly inhabitants. They wish to maintain their

could be used to determine his ID. The Use Case approach describes the essence of an action (indeed, in one version of uses case, Essential Use Cases [[3] L. L. Constantine, Essential modeling: Use Cases for user interfaces, *Interactions* **2** (2) (1995)., this property is emphasised in the name), and avoids describing detailed mechanisms to avoid committing the design to a particular approach prematurely.

independence and may wish never to interact deliberately and explicitly with a Smart Home in the way that users interact with more conventional software systems. If the elderly inhabitants are cognitively unimpaired but physically frail, it may be a matter of pride to achieve as much as possible without provoking intervention from the Smart Home. If their cognitive capabilities have begun to fail, they may simply be unaware that their Smart Home is ready and willing to help when they most need it to help them. Therefore, the majority of "interactions" between the inhabitant and the Smart Home may very well be unilateral interventions by the Smart Home with an inhabitant who is (deliberately or unintentionally) ignoring it. This fits poorly with the conventional Use Case formalism that requires a user goal to be expressed for each independent unit.

Smart Homes have secondary users such as the relative who installed it in order to assist the elderly inhabitant. The Smart Home's responsibility to the secondary users is to alert them when necessary and to collect messages from the Smart Home about interesting behaviors that the Smart Home decided not to react to. The goal that these users demand from the system is to monitor the inhabitant, assist, record, and react to whatever might happen. This also does not fit the classical definition of a user's goal in the Use Case terminology.

Although we may not be able to attribute a goal to the user in the interaction, it sometimes makes sense to attribute a goal to the system. Of course, if we suggest that, when an inhabitant leaves a pan of food cooking on the stove for an hour and the system turns the stove off, it is "trying" to prevent the house from burning down; we do not claim that the system has an explicit, conscious goal.

The second area of mismatch between conventional Use Cases and the Smart Home version concerns the structure of conventional interaction sequences. In most applications, interactions between the user and the system constitute a form of dialog (whether or not they involve the interaction components called dialog boxes). At the most primitive level, the user requests the application to perform a function and the application performs it. At more complex levels, the interactions may involve the transfer of multiple pieces of information, alternately generated by the user and the system. This alternating sequence of user output and system output is admirably captured by interaction sequences that comprise a numbered list, as is the earlier example. However, such a representation is not well suited to Smart Home "interactions" with elderly users, because they do not involve the same sort of give and take. For this reason, the Use Cases presented here include descriptions, but not interaction sequences. For a similar reason, Exception Scenarios are not always treated as part of the main Use Case; when the Smart Home takes an action to deal with an elderly dementia-sufferer's behavior, the action is not always consistent with the inhabitant's original goal. For example, when the inhabitant is discovered huddling in the bottom of the shower, having forgotten what they were doing there, and having forgotten how to get out of the shower, it is not appropriate to continue with their original goal of taking a shower. Instead, the elderly person needs to be warmed up, dressed, reassured, and checked for hypothermia. Consequently, in the list of Use Cases in the following section, most of the exception handling is separated out into different categories and different Use Cases from the ones that generate the exceptions.

3. SHMUC Single-Inhabitant Smart Home Use Cases

The Use Cases presented here have also been uploaded to the MUSE research group website at http://MUSE.massey.ac.nz/SHMUC. The Use Cases in the SHMUC repository are available as a resource for anyone in the Smart Environment community to use, and they are freely editable (although we reserve the right to moderate the edits if inappropriate material is uploaded).

In the Use Cases presented in this paper, there are four actors, three human, and one artificial. **Mary** is an elderly woman who is experiencing the early stages of dementia. Mary's daughter, **Debbie**, is unable to look after her mother because she lives elsewhere and has a full-time job and, in any case, wishes to respect her mother's strong desire to continue living as independent a life as possible for as long as possible. **Carita** is a professional carer who is on call to help Mary out of difficult situations, should this be necessary. The fourth actor, the artificial one, is **the Smart Home** (or **the system**) itself, an intelligent agent whose overarching goal is to support its elderly inhabitant in carrying on with a normal, safe and independent life, and to identify unusual behavior and act on it appropriately. Typical concerns would be to prevent the house from burning down, to ensure that Mary is going about her normal activities and to ensure that her blood sugar level is within acceptable limits. The Smart Home should complement existing facilities such as fire alarms or sprinklers. Overall, it should interfere in Mary's activities as little as possible, to remind Mary gently when things go wrong but she is expected to be able to correct the situation herself, to alert Debbie when something needs attention, but not immediately, and to request assistance from Carita when immediate intervention is required. It will be apparent that these responses have been arranged in increasing order of urgency. It is also strongly desirable that the minimum possible number of false positives and false negatives occur.

In the type of situation presented here (i.e., a single elderly inhabitant of a Smart Home), the terms *false positive* and *false negative* would normally indicate respectively that a problem had been detected where none exists, or that a problem that should have been detected and dealt with has gone unhandled. Note that this includes situations where the system reports a problem to Mary, but she fails to deal with it, a situation to which we shall return later. False negatives are problematic because they may put the inhabitant in danger. False positives are problematic because the human actors will eventually – or quickly – discount the system's alarms if it repeatedly "cries wolf." The three-level gradated response referred to above is intended to address both of these problems, by allowing small problems to be dealt with without demanding over-the-top external interference (a false positive) and by ensuring that information about major problems actually reaches someone who is competent and available to handle them. Note, however, that this gradation introduces another cause for both false positives and false negatives. If a problem is reported to someone whose interference is not required, then a false positive has effectively occurred, even though the problem was a real one, and needed to be attended to. If a problem is left for someone who is not competent or available to handle it, then a *de facto* false negative has occurred.

The situation in which we are proposing to deploy Use Cases is somewhat unusual, and we have introduced three new fields to fit them for the task. As the underlying purpose of the exercise is to assist us in developing techniques for detecting abnormal behavior, it is beneficial to have some idea of normal behavior in each case. For this reason, we include a field called *Norm*, which documents the inhabitant's normal behavior. A second new field, *Severity*, is used to capture the severity of the abnormal behavior documented in the Use Case. At this stage, the levels of severity are minimal, low, medium and high. Minimal severity abnormal behavior is interesting but not problematic, so the Smart Home can deal with it without making reference to any external authority. Low severity behavior prompts a warning to the elderly inhabitant of the Smart Home, to which she needs to respond. Medium severity alerts and low severity alerts to which the inhabitant has not responded are sent to the inhabitant's daughter. A high severity warning is sent to a carer who is available to attend the inhabitant of the Smart Home immediately.

Finally, each Use Case incorporates a field for discussion. This is called *System Design Implications*, as its purpose is to contain discussion of the conclusions that can be drawn about the design of the Smart Home, based on the Use Case.

Another unusual feature of the Use Cases presented here is that they have what might be called the Jeeves attribute, after the unobtrusive, ever-present, ever-helpful gentleman's gentleman in P.G. Wodehouse's Jeeves short stories (e.g. [20]). When circumstances are within normal boundaries, they specify no interference, and when their intervention is required, it is kept to a minimum

Although Use Cases are intended as a preliminary investigative tool, those presented here have already been classified into one of seven groups, according to the type of requirement they exemplify. These are illustrated in Fig. 1, which also shows that there are four major classes of Use Case: Use Cases that deal with abnormal behavior on the part of the inhabitant (which can be further broken down into spatial, temporal and pattern-of-action abnormalities); Use

Cases that deal with changes in the Smart Home environment; Use Cases that deal with changes in the context of the inhabitant's behavior; Use Cases that deal with an inadequate response to an earlier alert produced by the Smart Home.

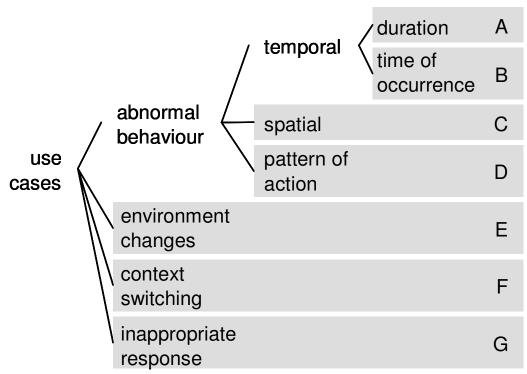

Fig. 1: The Use Cases presented in this paper can be grouped into seven categories

3.1. Use Cases Concerning Abnormal Duration (Class A)

The first broad category of Use Cases deals with situations in which an action takes more time or less time than expected. Many common activities normally take place within temporal limits. Most of us have at some time looked at the clock and said "Something's wrong." A family member may have returned from a six-hour trip after 20 minutes, or spent two hours in the cellar when their only reason for going down there was to fetch a hammer. The expected duration of such activities may be well defined or vague, but when we have decided that it is outside common limits, we start to consider raising the alarm. But there are decisions to be made. Are we being overly solicitous in raising the alarm already, or lax in leaving it so long? Should we mount a search ourselves or should we call in the police? Smart Homes have to make similar decisions. When is it appropriate to raise the alarm, and how strident an alarm should be raised? At whom should the alarm be directed?

Use Case A1: An Over-Long Shower

Goal
Mary wishes to take a shower.

Actors
Mary, the Smart Home, Debbie

Initial State
Mary was at home alone.

Scenario
Mary woke up at 8:00am as usual, and prepared to take a shower. She began her shower at 8:10am and at 8:40am she was still taking a shower, since the motion sensors in the shower room and the shower tap were on at that time. The system checked her ADL profile and found that she never normally took showers longer than 20 minutes, and that during winter they were generally even shorter. It alerted Mary to the excessive duration of her shower.

Norm
Mary's shower duration is normally between 10 and 20 minutes.

Severity
Low

Outcome
Mary, who had lost track of time, came out of the shower in response to the message, dried herself and dressed. A "no action required" message was sent to Debbie to keep her apprised of the state of her mother's behavior.

System Design Implications
Activities that take longer than they should may put the Smart Home inhabitant at risk. In this case, the activity took 30 minutes, 10 minutes longer than usual. This is not a large extension, in absolute terms, and it raises some further questions.

When does an activity such as a shower become longer than usual? The concept might be defined in absolute terms (5 or 10 minutes longer than average), relative terms (20%, 50%, 75% longer than average) or statistical terms (1, 2 or 3 standard deviations from the average). The last alternative seems most robust.

A second set of questions relates to classifying the situation. A longer-than-usual shower is interesting behavior because it is abnormal (outside the already known norm), but should it be treated as a problem? If so, what is its severity? And how does the system learn these things? Is it reasonable to expect a human – possibly Debbie, possibly Mary – to explicitly record a norm for each of Mary's ADLs?

Use Case A2: A Justifiably Short Shower

Goal
Mary wishes to take a shower

Actors
Mary, the Smart Home

Initial State
Mary was at home alone.

Scenario
Mary awoke on a very cold morning, so she took only 5 minutes in the shower, and then went into the living room and turned on the gas fire to warm up. The system detected from its knowledge base that the shower was much shorter than usual, but it was then able to correlate this with its knowledge of the weather and that Mary had used the heating.

Norm
Shower duration is normally 10 - 20 minutes

Severity
Minimal

Outcome
The system did not generate an alert message

System Design Implications
Though the duration of an activity is outside the norm, it may be justified for contextual reasons.

This Smart Home Behavior seems to rely on world knowledge. It is difficult to see how the Smart Home could recognize that information regarding the air temperature and Mary's use of the heating were relevant without some overall knowledge of how the world and people work. However, it could ask Mary why she had cut her shower short. To do so would imply that it had some way of allowing Mary or her carers to record reasons for behaviors.

Use Case A3: An Over-Long Nap

Goal
Mary wishes to take a nap

Actors
Mary, the Smart Home, Debbie

Initial State
Mary is at home alone.

Scenario
After lunch, Mary went into the living room to watch TV from 1:00pm to 1:30pm. At 1:30pm, she turned off the TV and then the system recognized that she was still on the sofa and so it assumed that she was taking a nap, as she often did. No activity was registered. At 3:30pm the system recognized that the nap was longer than usual. It reasoned that napping was not a dangerous activity and that Mary had had a disturbed night, and could therefore be tired. However, there was also the possibility that poor health was making Mary more tired than usual, or even that she was unconscious. The system therefore sent a message to Debbie to say that her mother was assumed to still be asleep, but the nap had lasted more than two hours.

Norm
The normal duration of Mary's afternoon nap is between 0 and 60 minutes

Severity
Medium

Outcome
The system sent an alert message to Debbie informing her that her mother was assumed to still be asleep, but that the nap had lasted more than two hours.

System Design Implications

A situation with an abnormal duration may be a problem if there is some danger (e.g., of health problems) involved.

The description above assumes a high degree of world knowledge on the part of the Smart Home. It is only justifiable to interpret the minimal information provided by simple Boolean sensor outputs in terms of such high-level concepts as "taking a nap" if those outputs are interpreted in the context of an ontology of world knowledge. It would be preferable if the system could make its decisions without needing to know about concepts such as "taking a nap."

3.2. Discussion of Abnormal Duration Use Cases (Class A)

It could be thought that once a behavior has been correctly identified, recognizing when its duration is too long or too short would be a simple task. However, the anomaly needs to be detected while the activity is still happening, meaning that not all of the cues to the behavior may have been seen. If we wait until the activity finishes, a dangerous situation may have occurred before the alert. Additionally, there are many contextual factors that affect the duration of an activity; Fig. 2 highlights some potential factors for the showering behaviors. Identifying, representing and accounting for them in the reasoning process is non-trivial.

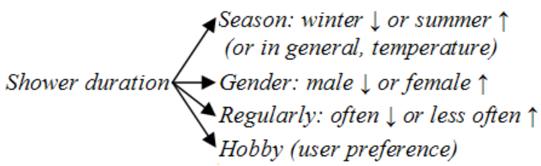

Fig. 2: Variations in showering behavior

Finally, it is not clear how much too long or short the duration of a behavior should be before it is considered abnormal. This obviously depends upon the statistical variance in the normal duration of the activity, and also what the behavior is – ten minutes too long in the oven is less serious than ten minutes too long in the microwave. This is a very difficult research problem that has been largely ignored so far. We believe that an ontology or similar representation of knowledge will be required.

3.3. Use Cases Concerning Time Of Occurrence (Class B)

The start time is a meaningful factor in Smart Home monitoring: an inappropriate activity start time could imply illness or even dementia. It is possible that a forgetful person can be reminded of activities that should have taken place.

Use Case B1: Variation in Shower Start Time

Goal

Mary wishes to take a shower

Actors

Mary, the Smart Home

Initial State

Mary is at home alone

Scenario

One morning, Mary awoke at 8:00am and it was very cold since winter was coming. So she decided not to take a shower immediately, intending to wait until 8:30am. The system checked that Mary did not take a shower from 8:00am to 8:20am as had occurred in the summer.

Norm

Mary's shower normally starts between 8:00am and 8:20am.

Severity

Low

Outcome

The system generated an alert to remind Mary to have a shower, who shrugged her shoulders and obediently went into the bathroom and showered.

System Design Implications

Mary's behavior was successfully, but probably unnecessarily, modified by the system. It is important for alerts – particularly the alerts to the inhabitant of the house – to be phrased very carefully so that there is a clear differentiation between alerts that carry advice and alerts that carry a safety-critical message.

Use Case B2: Taking Medicine after Midnight

Goal
Mary wishes to self-medicate for a stomach upset

Actors
Mary, the Smart Home, Debbie

Initial State
Mary is at home alone. She feels unwell.

Scenario
Mary awoke at 1:00am with a stomach upset and therefore decided to go and get some tablets from the medicine cabinet. The motion sensor outside the bedroom and the door sensor informed the system that she was leaving her bedroom, which was common for bathroom visits. However, when the sensor on the medicine cabinet fired, the system checked her medicine schedule and did not find any expected medication at that time.

Norm
Scheduled activities take place within a narrow band of times around the scheduled time.

Severity
Medium

Outcome
Since the system did not know what Mary was taking, an alert was sent to Debbie,

System Design Implications
Unusual start times for scheduled activities should generate a response.

Explicitly scheduled events are different from normal, regularly occurring, events such as showering. It can reasonably be assumed that they were scheduled because it is important that they occur at the scheduled time. Therefore, if they fail to occur at the scheduled times, or if they occur at other times when they are not scheduled to occur, then the behavior is both *interesting* and *problematic* and deserves to be drawn to the attention of a carer: in this case the severity of the behavior is *medium*, so it is Debbie who is alerted, and not Carita. To support such a distinction, explicitly scheduled events should have an *importance* index, so that an alert of the appropriate severity can be generated if the scheduled event does not occur.

Use Case B3: Late For Church

Goal
Mary wishes to go to Church

Actors
Mary, the Smart Home

Initial State
Mary is at home alone on Sunday before 9:00am

Scenario
Mary attends a church service that is held each Sunday morning at 9:30am. She usually leaves at 9:00am since the walk to church takes 20 to 25 minutes. One Sunday morning her friend rang and offered her a lift, and she accepted. The system noticed that she did not leave the house as usual at 9:00am

Norm
On Sunday mornings, Mary leaves the house to go to church within 5 minutes of 9:00am.

Severity
Low

Outcome
The system raised an alert for Mary.

System Design Implications
The Smart Home raised an unnecessary alert.

There is some data that the system cannot possibly know, and it will therefore reason incorrectly. When such situations occur, they need to be identified and corrected. If the Smart Home detects an anomalous start time and requests clarification from Mary, then she can explain that the behavior is unusual (i.e., the

Smart Home should not adjust its implicit schedule[2], but that it is acceptable in this instance.

Perhaps the first time this situation occurs, the Smart Home can also ask Debbie for clarification, and if she says that it is OK (essentially telling the system that her mother's word is to be relied upon), the Smart Home may only request clarification from Mary in future. Or maybe the Smart Home continues to ask Debbie for clarification, but less frequently.

Use Case B4: The Missed Visit

Goal
Mary wishes to alter her schedule without bothering her daughter

Actors
Mary, the Smart Home, Debbie

Initial State
Mary was at home alone on Monday afternoon.

Scenario
Mary always visits a friend at 3pm on Mondays. One day, she misses this regular visit, since she is feeling unwell. The system also detected that she opened the medicine cabinet at 2:30pm.

Norm
Mary leaves the house on Mondays; she does not normally open the medicine cabinet at 2:30pm.

Severity
Medium

Outcome
Based on the two incidents – Mary staying at home when she would normally go out, and opening the door to the medicine cabinet – the system sent an alert to Debbie.

System Design Implications
It should be possible to merge data from different sources to identify an abnormal situation. In this case, the opening of the medicine cabinet indicates that Mary is unwell, and that events in her normal schedule may be disrupted if they are inessential. This has two implications for the system's world knowledge. It needs to know whether opening the medicine cabinet is an event that can justifiably disrupt the ordinary schedule, and it needs to know that Mary's visit to her friend (which, of course, it only knows about at the semantically poor level of *Mary leaves the house on Mondays*) is inessential.

3.4. Discussion of Abnormal Start Time Use Cases (Class B)

As with the abnormal duration, the abnormal start times presented in the three preceding Use Cases seem easy to detect. However, there are contextual and other issues that affect things:

i) While some people have regular schedules and fixed activity start times, others are more variable in their timings. There are also some behaviors that, even amongst the most dependable people, move from their normal times due to some sudden reasons such as in Use Case B2 (Taking Medicine after Midnight). For this problem, it may be that behaviors will need to be categorised into regular and irregular ones. However, the problem in Use Case B2 is still hard to avoid and we almost have to accept it.

ii) The variation in activity start time is affected not only by inhabitant's preference but also by the other contextual factors. Identifying those factors may help to increase the detection accuracy. For example, without context information, the start time of Mary's shower should be from 8:00am to 8:30am. However, it would be more accurate if the system could distinguish between summer time (8:00am to 8:20am) and winter time (8:30am, which may be

[2] Here, we use the phrase *implicit schedule* to refer to a schedule that the Smart Home derives from observation of the inhabitant's behavior. This is in contrast to an *explicit schedule* which would be created and maintained by explicit interactions by the inhabitant or a carer.

updated as more situations are observed); this requires the system to be able to consider contextual information.

iii) As with duration, there is a question about how long the system should wait before issuing an alert. Related to this is the question about what is a suitable output. For example, if the system could interact with Mary instead of raising an alarm, it could ask if she had forgotten to go to church, or issue other reminders.

3.5. Spatially Abnormal Behaviors (Class C)

Performing an action in the wrong place may endanger people or signify that something has gone wrong. The following scenarios will discuss how the spatial properties of an activity can help to identify abnormal behavior.

Use Case C1: Lying Down in the Kitchen

Goal
Mary wishes to prepare breakfast

Actors
Mary, the Smart Home, Carita

Initial State
Mary was at home alone.

Scenario
At 8:15am, Mary went to the kitchen to prepare breakfast. She put some bread on a plate and then lay down on the floor. The behavior was recognized, and its spatial properties checked. As this behavior should not be seen in the kitchen it was potentially serious.

Norm
Inhabitants do not normally lie down in the kitchen.

Severity
High

Outcome
After querying Mary about her behavior and receiving no response, the system sent an alarm to Carita

System Design Implications
Some behaviors should generate an immediate reaction, as they are potentially very significant.

There is a class of activities that fall outside the bounds of normal behavior and can be *prima facie* assumed to be both *interesting* and *problematic*. This should reduce the computational effort involved in deciding how to react to an observed activity. However, complementing that is the difficulty of foreseeing all possible inappropriate behaviors. As mentioned earlier, it is difficult enough to build a world model that allows for normal behaviors, but the size of the problem is potentially much larger if the system has to detect and classify all possible dangerous abnormal behaviors.

3.6. Discussion of Spatially Abnormal Use Cases (Class C)

Detecting abnormality of spatial activities requires the spatial data to be stored by the system and attached to a behavior. In comparison with the previous Use Cases, the Use Cases in this category are more static, as the information does not change frequently. A further issue that may arise is that if the system had sufficient sensors to recognize where somebody was lying, and information about how they got there (carefully or abruptly) then it could better identify the need for an urgent response to a fall. Without the latter information, the best that it can do is to recognize a behavior that is seen frequently but where the location is unexpected.

3.7. Abnormalities in Patterns of Behavior (Class D)

One of the major challenges of identifying behavior patterns is that their appearance in terms of sensor patterns can vary immensely between people, and even within the activities of one person. For example, there are many different ways to cook dinner, depending upon what it is; evidence from various Smart Home datasets suggests that between 4 and 58 actions are needed depending upon the type of food prepared [15], and other behaviors exhibit similar variation. This is an area where it can be particularly difficult to decide what a Smart Home should be able to detect, and how to reduce the risk of false positives. We use three scenarios for the common and easily understood task of making a cup of tea to illustrate the complexity of identifying errors even in simple task.

Use Case D1: Making Tea with Sugar

Goal
Mary wishes to make a cup of tea

Actors
Mary, the Smart Home, Debbie

Initial State
Mary is at home alone.

Scenario
During the training phase, the system identified a tea-making behavior with the "syntax" shown in Fig. 3.

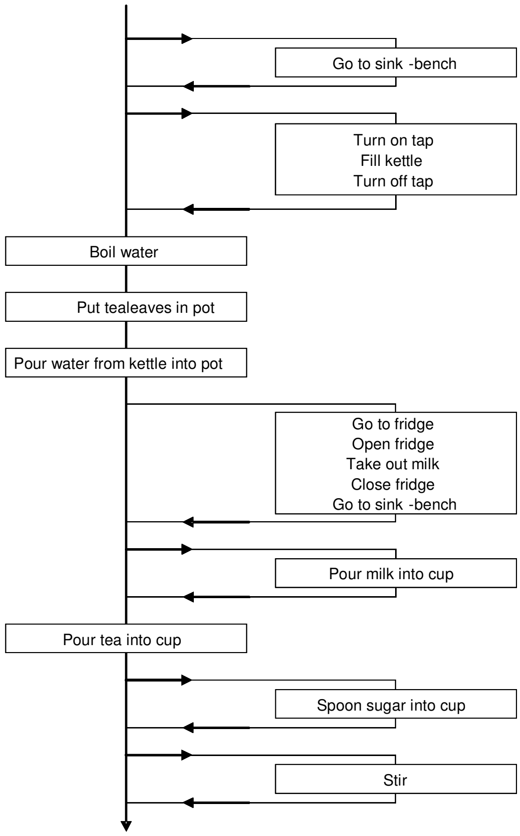

Fig. 3: Inferred "syntax" for making a cup of tea

After training, the system was ready to monitor Mary's activity. One afternoon Mary made a cup of tea without milk, but with sugar, using the following set of actions:

go to sink-bench
turn on tap, fill kettle, turn off tap
put tealeaves into pot
boil water
pour water from kettle into pot
pour tea into cup
spoon sugar into cup
stir

This sequence does not match the learned pattern exactly, as the items in bold are reversed in order with respect to the syntax created during training, so the sequence was considered as a novelty by the system. However, it did not cause an immediate alarm, as the system identified that the order of two actions, i.e. **put tealeaves into pot** and **boil water**, does not affect the final state of the activity. Therefore, the system did not create an alert, but modified its representation of tea making instead.

Norm
The tea-making sequence conforms to the syntax specified by the Finite State Machine.

Severity
Minimal

Outcome
The activity pattern was automatically updated, and a "no action required" notification was sent to Debbie.

System Design Implications
This Use Case deals with multiple valid activity orderings.

Activities often comprise a partially ordered sequence, and there is no guarantee that observation of any number of instances will reveal all the orderings. It is therefore important that the system should incorporate a mechanism to recognize that the events that make up an activity have occurred out of the normal sequence - or equivalently, that an order of events that resembles, but does not correspond exactly to any previously observed order, may be a previously unseen but valid order.

The system could not infer this without external input from a competent source (which might rule out an inhabitant with dementia).

Use Case D2: Making Tea with Cold Water

Goal
Mary wishes to make a cup of tea
Actors
Mary, the Smart Home, Debbie

Initial State
Mary was at home alone.

Scenario
Mary went to the kitchen to make a cup of tea and this sequence of sensor observations was detected:

go to sink-bench
turn on tap
fill kettle
turn off tap
boil water
put tealeaves in pot
t**urn on tap**
fill pot with water
pour tea into cup
stir

Based on this information, the system needed to identify the behavior. Although it does not directly match any of the stored behaviors, the closest match is to the tea making activity.

Norm
The tea-making sequence conforms to the syntax specified by the tea-making Finite State Machine in Use Case D1.

Severity
Minimal

Outcome
The system reports an abnormal behavior in a "no action required" message to Debbie.

System Design Implications
In this Use Case, an essential property of the tea-making behavior has been replaced by something inappropriate. The additional behavior is *interesting* in that it does not conform to the norm, and probably *problematic*, in that it may indicate that Mary is confused. While this classification (probably problematic) is an easy one for us to reach, based on our world-knowledge appreciation that a cup of tea made with cold water will taste particularly foul, and that anyone who makes tea that way may be confused, the system has no such model. How could it come to a similar conclusion?

Use Case D3: Making Tea with Hot and Cold Water

Goal
Mary wishes to make a cup of tea

Actors
Mary, the Smart Home, Debbie

Initial State
Mary was at home alone

Scenario
In the afternoon, Mary made another cup of tea and the following sensor sequence was observed:

go to sink-bench
turn on tap, fill kettle, turn off tap
boil water
put tea in pot
pour water from kettle into pot
pour tea into cup
turn on tap, put cold water into cup
spoon sugar into cup
stir

This sequence was classified as unusual since the action “put cold water into cup” was not in the learned model of tea making.

Norm
The tea-making sequence conforms to the syntax specified by the Finite State Machine.

Severity
Minimal

Outcome
A warning message was recorded in the system as above. When Debbie later reviewed the notifications from the system, she considered this sequence, and realized that it was perfectly reasonable, if Mary wished to drink the tea immediately, and had added cold water to cool it down.

System Design Implications
This is another situation in which the system needs external input to determine whether or not the inhabitant's behavior is reasonable. As in case D1, it is reasonable to query the inhabitant about this.

Use Case D4: Taking a Shower while Cooking

Goal
Mary wishes to prepare lunch

Actors
Mary, the Smart Home, Debbie

Initial State
Mary is at home alone.

Scenario
Mary began to prepare her lunch. Her activities are shown in Fig. 4:

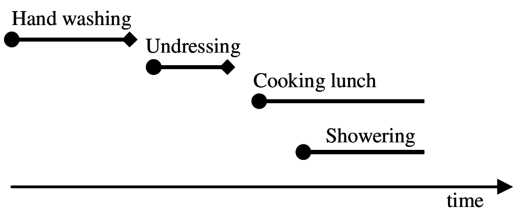

Fig. 4: Mary's sequence of actions while preparing lunch

Norm
The inhabitant of a Smart Home does not normally mix showering and cooking

Severity
Low, changing to medium

Outcome
The system queried Mary about her behavior, and when it did not receive a reply, sent an alert to Debbie.

System Design Implications
Some pairs of behaviors may sensibly be intermixed and some may not. For example, it is safe to have a cup of tea while cooking a meal but not to have a snooze while a pan of chips is deep-frying in oil. These examples clearly exemplify safe and unsafe behavior pairs, but there is a gradation between these

easily classifiable extremes, not a sharp cutoff. For example, if an inhabitant set the oven to cook a roast for two hours and, feeling well satisfied, went for a nap, setting the bedside alarm to ring half an hour before the roast was ready, she or he might well feel aggrieved if the Smart Home raised a high-severity alarm. But what if an inhabitant with highly reliable sleeping patterns dispensed with the alarm? What, then, if the inhabitant were boiling a pot of potatoes? The smell would be awful, the pot would be ruined, but the house would probably not burn down.

When is a pair of mixed behaviors safe and when is it unsafe? From a risk management perspective, it is tempting to treat activities with any element of peril as unsafe, but from the perspective of gaining and maintaining the inhabitant's cooperation with the Smart Home, it is desirable to minimize the extent to which it interferes with their way of life.

3.8. Discussion of Patterns of Behavior Use Cases (Class D)

These Use Cases demonstrate two difficult aspects of behavior recognition, that there may be wide variation in behavior presentation and that the difference between safe and unsafe behaviors can be subtle. Use Cases D1-D3 suggest that the system should be able to identify trivial changes to a learned behavior and modify its representation unilaterally. This is a significant challenge. It would need to be able to infer, from its sensor observations, that the essential features of the syntax for tea-making include putting tea into a pot and adding hot water, and that milk and sugar are optional, and they can be added to the cup at any stage without affecting the outcome. It may be possible to simplify this problem by focusing upon the final state of the sequence rather than tracing the whole progress.

3.9. Use Cases Concerning Changes in the Smart Home Environment (Class E)

The Smart Home environment is dynamic; it can change to meet any new requirements of the inhabitant, and equipment and fittings within the home can break and subsequently be replaced or repaired. We want to be able to add new things into the house without the Smart Home system having to be retrained from the beginning. The question of what sensors are available and attached to objects is not discussed in this paper, but it is another part of this problem.

Use Case E1: New Blender Installed

Goal
Debbie's goal - to give her mother a new, safer blender

Actors
Mary, the Smart Home, Debbie

Initial State
There is a blender in Mary’s house.

Scenario
Mary sometimes forgets to turn appliances off after using them, which can be dangerous. For example, she often forgets to turn the blender off.

One day, Debbie discovered a blender that turns itself off every 90 seconds, so that it cannot be accidentally left on. She purchased this device for her mother, and replaced the old one. Any behaviors that involved the blender will change, but they can still be based on the experience that was acquired from the old blender until additional evidence is seen.

Norm
Some devices (e.g. Mary's first blender) are normally turned off after use; some, (e.g. her new one) are not.

Severity
Minimal

Outcome
The behaviors that involved the blender were progressively updated as evidence accumulated.

System Design Implications
It is important to minimize the extra learning required when a new device is installed.

Appliances can be categorized into those, such as blenders, hair-dryers, and stoves, that require turning off after use, and those that do not require turning off, either because they turn themselves off (Mary's new blender, modern kettles, toasters) or because they are left turned on all the time (refrigerators). In the database of world knowledge used by the Smart Home, each appliance would have an operation profile that the system would learn over a period. It is possible that, as smart environments become more common, manufacturers may equip appliances with operation profiles that specify their command vocabulary and network ports so that they can be controlled using voice commands via a standardized Smart Home interface. However, further investigation of such developments is outside the scope of the current discussion.

3.10. Discussion of Use Cases Concerning Changes in the Smart Home Environment (Class) E

When new devices or sensors are installed into the Smart Home, the system should be able to learn about them. This can be done in various ways, from starting again with any affected behaviors, through allowing them to be progressively modified as more evidence of the changed behavior is acquired, to the use of an ontology system to identify generalisations of the modification and make intelligent deductions about the system.

3.11. Use Cases in which ADL Structure Changes (Class F)

The system should be able to adapt to changes in the high-level structure of an inhabitant's ADLs. For example, the structure of Mary's day may follow one pattern on weekdays and change to quite another during the weekend. Or she may – as in the Use Case below – develop a new interest that changes her standard pattern of behavior.

Use Case F1: The Working Week

Goal
Mary wishes to broaden her horizons by taking on a volunteer position at an ESOL (English as a Second Language) Centre

Actors
Mary, the Smart Home

Initial State
Everything is normal.

Scenario
Bored with being at home alone, Mary registered for a volunteer position at an ESOL Centre to help foreigners learn English. Each level of the course takes about 4 weeks but there is no fixed schedule. Whenever the Centre has enough learners, they call Mary in the morning, and she assists with teaching at the Centre instead of watching TV, which she would normally do in the morning.

On the first morning of this new activity, Mary prepared to go to the ESOL Centre for a new class. When she was going out, the system checked her behavior set and recognized that this was a strange activity, since she was normally at home in the morning. However, it did not send a message to Debbie, but checked other ADL patterns to find an appropriate one for her. It found that there was a pattern in which Mary does go out in the morning if she has a class. This pattern was created when Mary registered for a volunteer position at the centre. So this pattern was used for monitoring her activities during that day.

Norm
Mary normally watches TV at home in the mornings

Severity
Minimal, because the system found an explanation for what would otherwise have been a medium-severity behavioral abnormality.

Outcome
The system recognizes that these behaviors are normal for this pattern.

System Design Implications
The system should be able to recognize and distinguish between behavior sequences that belong to one pattern and behavior sequences that belong to an alternative pattern, and treat them as mutually exclusive.

In this particular case, the context was created when Mary registered for the volunteer ESOL position. Was it something that she created explicitly herself? Did her daughter or Carita, having discussed her change of activity, create it? Or did the system learn it the first time it happened?

3.12. Discussion of Use Cases in which ADL Structure Changes (Class F)

Detecting and distinguishing between patterns of behavior is an interesting aspect of the Smart Home application, especially for human activity recognition and abnormal behavior detection, since both depend strongly upon the context. It represents the system's ability to adapt to changes in context, and therefore plays an important role in the accuracy of recognition and detection. Moreover, switching between states raises the idea of building a ubiquitous Smart Home system that can aggregate data from many other sources, e.g., from many different Smart Homes, and use it inheritably. For example, suppose that each house could inform others when their inhabitant was ill. This could lead to an early warning system for pandemics, where other houses warn their inhabitants to wash their hands frequently and stock up on food in case they become ill.

3.13. Use Cases concerning Inappropriate Response (Class G)

The concept of severity has been introduced into the Use Cases presented here. This makes it easy to determine whom the system response should be targeted at, but it introduces subsidiary problems. First, as has been noted, if the severity is incorrect, then someone is going to be bothered unnecessarily, or left unaware of a developing problem that they should be dealing with. Secondly, if an alert at a particular level of severity is unacknowledged, it needs to be escalated so that the carer at the next highest level can deal with it.

Use Case G1: Inadequate Response

Goal
The system wishes to alert Mary to the fact that her shower has lasted an unusually long time.

Actors
Mary, the Smart Home, Debbie

Initial State
Mary has been in the shower for longer than the expected duration of a shower.

Scenario
Mary has become confused, and doesn't know what she's doing in the shower; indeed she's forgotten what a shower is, or how to open the door. She has been alerted by the system to the unusual length of her shower, but does not know how to deal with the alert.

Norm
Mary responds to the low-severity alerts that are directed at her.

Severity
Low (which turns out to be incorrect)

Outcome
The hot water runs out, and Debbie finds Mary two hours later, huddled in the corner of the shower and shivering violently. She is admitted to hospital suffering from hypothermia.

System Design Implications
When no response to an alert occurs within a certain time, the system should, in the first instance, generate a higher-severity alert. In this case, increasing the severity by one step (from low to medium) is not enough because medium severity alerts are sent to Debbie, who does not normally respond immediately. Perhaps it should be possible to set a severity–if-not-responded-to property for each abnormal behavior (which would normally default to the next level up). In the longer term the system needs upgrading to allow for the more severe nature of Mary's dementia. Should this be a standard system function, or should it require explicit reprogramming?

Use Case G2: Correction of a Rule

Goal
Mary's wishes to take a shower

Actors
Mary, the Smart Home, Debbie

Initial State
Mary is at home alone

Scenario
One morning, Mary awoke at 8:00am and it was very cold, since winter was coming. So she decided not to take a shower immediately, intending to wait until 8:30am. The system checked that Mary did not take a shower from 8:00am to 8:20am as had occurred in the summer. Therefore it generated a low-severity alert to remind Mary to have a shower. When Mary did not heed the alert, the system upgraded it to medium severity and sent a warning message to her daughter, Debbie. However, although Mary did not take the shower when she got the alarm, she did take the shower – just 10 minutes later, at 8:30am, as she had intended. After work, Debbie found the message and on checking the system records, recognized that this was an incorrect inference that had occurred because the system had not observed this winter-time variation of her mother's showering activity. Debbie then provided feedback to the system to update this activity start time.

Norm
Mary's shower normally starts between 8:00am and 8:20am.

Severity
Low

Outcome
The rule about the start time was updated

System Design Implications
There may be an acceptable variation in the start time of an activity.

However, in general, it is probably safe to assume that an activity start time that is outside the norm (say, 2 standard deviations from the mean) is an *interesting* but not an inherently *problematic* behavior. Therefore it is acceptable to request external (human) input regarding the classification of the behavior, and it may not be necessary for the Smart Home to rely on pre-loaded world knowledge.

Use Case G3: Unacknowledged Low-severity alert

Goal
System wishes to alert Mary to the fact that her shower has lasted an unusually long time.

Actors
Mary, the Smart Home, Debbie

Initial State
Mary has been in the shower for 31 minutes, the Smart Home has issued a low-severity alert to Mary, suggesting that it is time she exited from the shower.

Scenario
Mary was perfectly fine, but, with no particular schedule for the day, was luxuriating in the feeling of hot water running over her body, and stayed in the shower for an unusually long time. When the Smart Home issued an alert after 31 minutes of showering, she stayed in the shower a minute longer, but then turned off the water and got out, dried off and dressed herself. In the meantime, the Smart Home, having been upgraded as suggested in Use Case G1, had issued an alert to Mary's daughter Debbie.

Norm
Mary responds to alert messages from the Smart Home

Outcome
Debbie receives an unnecessary alert and is annoyed by the behavior of the Smart Home.

System Design Implications
There needs to be a way of recognizing that, although Mary's behavior is outside the accepted limit of normality (31 minutes in the shower), it is not dangerous, and the alert that was issued after a 1-

minute extension should be cancelled. This suggests that for any behavior with an allowable range of durations, there should be a range of values that generate an alert, but that ceasing the behavior within some extended time should cause the alert to be cancelled. It also suggests that the inhabitant should be able to extend their behavior for a certain amount of time, when the original alert arrives ("I'll just have another five minutes").

Use Case G4: Ambiguous Sensor Output

Goal
Mary just wants to be tidy

Actors
Mary, the Smart Home, Debbie

Initial State
A bottle of antacid is sitting on the vanity in the bathroom

Scenario
While tidying up on Monday afternoon, Mary notices that she forgot to replace the bottle of antacid that she took out of the medicine cabinet the previous evening. She opens the door and replaces it

Norm
The medicine cabinet door is not normally opened in the middle of the afternoon, and Mary is not scheduled to take any medication till after dinner on Monday

Severity
Medium

Outcome
The Smart Home treats unscheduled medication as a medium-severity abnormal event, and it sends an alert to Debbie.

System Design Implications
The reliability of inferences made by the Smart Home from low-level sensor inputs depends on the number of inputs; when the number is large, inference that the inhabitant is actually carrying out the inferred behavior can be made with a high degree of confidence; when the number of inputs is small the confidence is correspondingly smaller.

It is probably not possible to disambiguate single-sensor inputs without world-knowledge or external input. And each time the system makes an incorrect inference and sends an unnecessary alert, confidence in the inferences is eroded. This is a significant problem. The Smart Home could check with Mary to get her to explain what she's doing, but it is not easy to see how Mary could respond in general to a query about whether or not her action is reasonable – especially if her actions are the result of diminished cognitive abilities.

Use Case G5: Correcting an Inappropriate Response

Goal
Debbie wishes to reclassify a behavior from abnormal to normal.

Actors
Debbie, the Smart Home, Mary

Initial State
A previous behavior was incorrectly classified as abnormal

Scenario
When making a cup of tea, Mary added a little cold water to moderate the temperature; the Smart Home had not encountered this behavior before and recorded a warning message. When Debbie found this, she decided that the behavior needed to be reclassified

Norm
Tea-making with added cold water is abnormal

Severity
Not Applicable

Outcome
The syntax for tea-making is updated

System Design Implications

Although the syntax for ADLs may vary widely between one person and another, individuals tend to have "relatively fixed routines for making tea in their own kitchen" [12]. However, that does not mean that an individual's routine never varies and the system therefore needs to incorporate a mechanism for easy updating of ADL syntax. This might prove a difficult problem; the syntax shown in Fig. 3 for making a cup of tea is very simple, and a graphical representation was chosen to make it easy to understand, but many people would find it difficult to create or modify such a syntax.

4. Conclusions

We have proposed Use Cases as a requirements-analysis tool in the context of Smart Homes, and presented a set of sample Use Cases, focusing mainly on abnormalities in ADLs, and the implications that these have for the design of the Smart Home. There are other areas that could also have been explored. In particular, the approach to Smart Home design that is espoused by the MUSE research group involves a preliminary training phase, and continuous upgrading of the in-use system. Neither of these areas has been explored, although upgrading has been touched on briefly, notably in Use Case category G. A web repository of Use Cases has been established at http:// MUSE.massey.ac.nz/SHMUC. The website has been made publicly available as a resource which all members of the Smart Environment community may draw on, and contribute to. It has been constructed using the Semantic Drilldown extension [10] to Semantic MediaWiki (the technology underlying Wikipedia) [13]. It has facilities for adding to and editing the set of Use Cases, for supporting discussions and for maintaining a history of changes. The Semantic Drilldown extension supports the hierarchy of Use Case categories that was shown in Fig. 1.

The field that is investigated in this work – a Smart Home for elder-care – is an unusual application area for Use Cases, in that the majority of "interactions" between the Smart Home and its inhabitant may very well be unilateral interventions by the Smart Home in the ADLs of an inhabitant who is deliberately or unintentionally ignoring it. This contrasts with conventional software systems in which the user deliberately initiates the interaction in pursuit of some goal. Consequently, we have found it useful to modify the structure of the Use Cases to include fields called *Norm* – which documents the inhabitant's normal behavior – *Severity* – which is used to capture the degree of danger associated with the abnormal behavior documented in the Use Case – and *System Design Implications* – which is used to contain discussion about the significance of the Use Case for the design of the system as a whole.

Analysis of the Use Cases produced a preliminary taxonomy of Use Cases for a Smart Home with a single elderly inhabitant. In this taxonomy, which was pictured in Fig. 1, there are four major classes: Use Cases that deal with abnormal behavior on the part of the inhabitant (which can be further broken down into spatial, temporal and pattern-of-action abnormalities); Use Cases that deal with changes in the Smart Home environment; Use Cases that deal with changes in the context of the inhabitant's behavior; Use Cases that deal with an inappropriate response to an earlier alert produced by the Smart Home. At the leaf node of the taxonomy there are seven categories. It should be emphasized that the taxonomy is tentative and incomplete.

4.1. Conclusions from the Use Cases

In the process of creating and reviewing the Use Cases, a number of design decisions have emerged. Some of the design decisions documented below are generalizations, amalgamations and extensions of decisions documented in the individual Use Cases.

- Some behaviors should generate an immediate reaction, as they are potentially very significant.
- When no response to an alert occurs within a certain time, the system should generate an alert of the next highest severity.
- Acceptable start times and durations for ADLs should be ranges and not single values
- An ADL with a numeric value (duration, start time, number of repetitions of an activity) is abnormal (*interesting)* if it deviates more than 2 standard deviations

from the mean, but there may be contextual reasons that prevent it from being *problematic*.

- There needs to be a way of recognizing that, for some numeric values, a value slightly outside the accepted limit of normality is not dangerous, and that an alert that was issued may be cancelled. This suggests that for any behavior with an allowable range of durations, there should be a range of values that generate an alert, but that ceasing the behavior within some extended time should cause the alert to be cancelled.
- An importance index should be associated with scheduled events so that the severity of missing a scheduled event can be reliably assessed, and unnecessary alerts can be avoided.
- The Smart Home should be able to ask its inhabitant for a reason for abnormal behaviors and avoid raising unnecessary alerts.
- Reasons given by the inhabitant for deviations from abnormal behavior should be recorded and reviewed by a carer to check for problems.
- In the database of world knowledge used by the Smart Home, whether or not an appliance needs to be turned off would be a Boolean property of each appliance. In general, an operation profile needs to be maintained for each appliance.
- The system should be able to recognize and distinguish between activities that belong to one pattern of activities and activities that belong to an alternative pattern, and treat them as mutually exclusive.

A number of Use Cases gave rise to possible design decisions.

- The Smart Home could ask the inhabitant the reasons for low-severity abnormal events. To do so would imply that it had some low-impact way of allowing an inhabitant to record reasons for behaviors.
- It should be possible to merge data from different sources to identify an abnormal situation. (Could this be a high-level application of Hidden Markov Models?)
- It might be possible to build a ubiquitous Smart Home system that can aggregate data from many other sources, e.g., from many different Smart Homes, and use it inheritably.
- It may be that behaviors will need to be categorized into *regular* and *irregular*, so that irregular behaviors do not trigger alerts even when they occur at odd times.
- When a parameter for an ADL is updated, it may be possible to use the ontology of world-knowledge to generalize it and propagate the change to other ADLs.

A number of challenges have also emerged.

- Several of the Use Cases assume a high degree of world knowledge on the part of the Smart Home. It is worth searching for techniques to avoid relying on world knowledge. For example, the idea of recording low-severity alerts that were handled by the inhabitant, so that a carer can review them and detect subtle patterns of degeneration in the inhabitant's mental state reduces the need for intelligence on the part of the Smart Home.
- Anomalies in duration or start time may need to be detected before the activity is complete, so that interventions occur in time in dangerous situations such as the inhabitant getting stuck in the shower.
- There is some data that the system cannot possibly know, and it will therefore reason incorrectly. When such situations occur, they need to be identified and corrected.
- Without information about how a situation such as lying on the floor came about, the best that the system can do is to recognize a behavior that is seen frequently but where the location is unexpected.
- From a risk management perspective, it would probably seem appropriate to treat anything that had any element of peril as unsafe, but from the perspective of gaining and maintaining the inhabitant's cooperation with the Smart Home, it would be

appropriate to minimize the extent to which it interfered with their way of life. It is desirable to use techniques for disambiguating such situations, such as asking the inhabitant; recording low-severity abnormalities for a carer to oversee, rather than issuing an immediate alert.

5. Future Work

The Use Cases described here were focused on detecting abnormalities in the inhabitant's Activities of Daily Living. Other sets of Use Cases need to be added to these. For example, Use Cases are needed for Smart Home training activities and for updating the syntax of existing ADL descriptions.

References

[1] R. S. Bucks, Ashworth, D. L., Wilcock, G. K., Siegfried, K., Assessment of Activities of Daily Living in Dementia: Development of the Bristol Activites of Daily Living Scale, *Age and Ageing* **2,** (1996), 113-120.

[2] A. Cockburn, *Writing Effective Use Cases*. (Addison-Wesley, 2001).

[3] L. L. Constantine, Essential modeling: Use Cases for user interfaces, *Interactions,* **2** (2) (1995).

[4] H. Guesgen, W. and S. Marsland, in: *Handbook of Research on Ambient Intelligence and Smart Environments: Trends and Perspectives*, edited by F. Mastrogiovanni (2010).

[5] H. Guesgen, W. and S. Marsland, in: *Handbook of Ambient Intelligence and Smart Environments*, edited by H. Nakashima, Aghajan, H., Augusto C. J. (Springer, Berlin, Germany, 2010), pp. 609-634.

[6] H. Hagras, V. Callaghan, M. Colley, G. Clarke, A. Pounds-Cornish and H. Duman, Creating an ambient-intelligence environment using embedded agents, *IEEE Intelligent Systems* (2004), 12-20.

[7] S. Helal, W. Mann, H. El-Zabadani, J. King, Y. Kaddoura and E. Jansen, The Gator Tech Smart House: a programmable pervasive space, *Computer,* **38** (3) (2005), 50-60.

[8] I. Jacobson, Object-Oriented Software Engineering - a Use Case Driven Approach, presented at: TOOLS (10) 1993 (unpublished).

[9] C. D. Kidd, R. Orr, G. D. Abowd, Atkeson, C. G., I. A. Essa, B. McIntyre, M. E. D., T. Starner and N. W., The aware home: A living laboratory for ubiquitous computing research, *CoBuild'99 Prodeedings of the Second International Workshop on Cooperative Buildings, Integration Information, Organization, and Architecture* (1999), 191-198.

[10] Y. Koren, Extension:Semantic Drilldown, 2009, *http://www .mediawiki.org/wiki/Extension:Semantic_Drilldown*, accessed 4 April, 2010.

[11] N. Robert, Taewoon, K., Mitchell, L., Stephen, K., Living quarters and unmet need for personal care assistance among adults with disabilities, *Journal of Gerontology: Social Science* (2005).

[12] J. Rusted and L. Sheppard, Action-based Memory in Alzheimer's Disease: a Longitudinal Look at Tea Making, *Neurocase,* **8** (2002), 111-126.

[13] Semantic_MediaWiki, Semantic MediWiki, *http://semantic-mediawiki.org/wiki/Semantic_MediaWiki*, accessed 28 April 2010.

[14] H. Sharp, J. Preece and Y. Rogers, *Interaction Design - Beyond Human Computer Interaction*. (John Wiley and Sons, 2007).

[15] E. M. Tapia, Activity Recognition in the Home Setting Using Simple and Ubiquitous Sensors, *http://courses.media. mit.edu/2004fall/mas622j/04.projects/home/*, accessed 28 October 28, 2009.

[16] E. M. Tapia, S. S. Intille and K. Larsson, in *Proceedings of PERVASIVE 2004*, edited by A. Ferscha and F. Mattern (Springer-Verlag, 2004), pp. 158-175.

[17] United Nations Department of Economic and Social Affairs, World Population Prospects: The 2006 Revision, (2006).

[18] University of Colorado, The Adaptive House, 2009, *{http://www.cs.colorado.edu/~mozer/house/*, accessed 28 October 2009.

[19] W3C, D. Chamberlin, P. Fankhauser, D. Florescu, M. Marchiori and J. Robie, XML Query Use Cases, *http://www.w3.org/TR/xquery-use-cases/*, accessed 3 December, 2009.

[20] P. G. Wodehouse, *My Man Jeeves*. (Bed Books, 2005).

[21] G. M. Youngblood and D. J. Cook, Data mining for hierarchical model creation, *IEEE Transactions on Systems, Man, and Cybernetics, Part C,* **37** (4) (2007), 561-572.