

Spatio-Temporal Reasoning and Context Awareness

Hans W. Guesgen and Stephen Marsland

1 Introduction

Smart homes provide many research challenges, but some of the most interesting ones are in dealing with data that monitors human behaviour and that is inherently both spatial and temporal in nature. This means that context becomes all important: a person lying down in front of the fireplace could be perfectly normal behaviour if it was cold and the fire was on, but otherwise it is unusual. In this example, the context can include temporal resolution on various scales (it is winter and therefore probably cold, it is not nighttime when the person would be expected to be in bed rather than the living room) as well as spatial (the person is lying in front of the fireplace) and extra information such as whether or not the fire is lit. It could also include information about how they reached their current situation: if they went from standing to lying very suddenly there would be rather more cause for concern than if they first knelt down and then lowered themselves onto the floor. Representing all of these different temporal and spatial aspects together is a major challenge for smart home research. In this chapter we will provide an overview of some of the methodologies that can be used to deal with these problems. We will also outline our own research agenda in the Massey University Smart Environments (MUSE) group.

Our interpretation of the smart home can be seen in Figure 1. We assume that unspecified sensors make observations of the house, and that changes in the observations, and hence the output, of these sensors in some way represent the behaviour of the human inhabitants. We are not directly interested in the signal processing that

Hans W. Guesgen
School of Engineering and Advanced Technology, Massey University, Palmerston North, New Zealand, e-mail: h.w.guesgen@massey.ac.nz

Stephen Marsland
School of Engineering and Advanced Technology, Massey University, Palmerston North, New Zealand, e-mail: s.r.marsland@massey.ac.nz

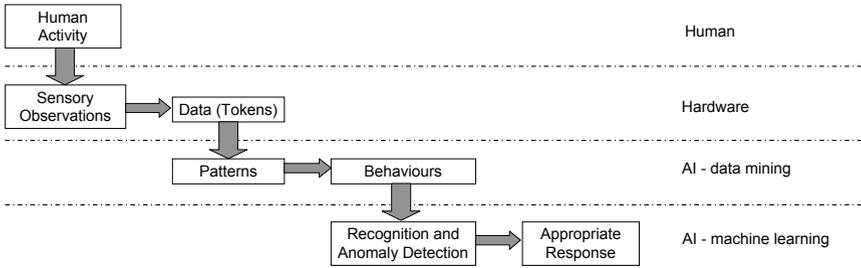


Fig. 1 The smart home problem consists of interpreting sensory observations caused by human behaviour in order to recognise human activity and respond appropriately.

is inherent in interpreting sensory output; rather we assume that these signals are presented to the smart home controller in the form of ‘tokens’. These could be a direct representation of current sensor states, polled at regular time intervals, or indications that something has changed (in the former case the output could be ‘switch is off’, ‘switch is off’, ‘switch is off’, and so on, while in the latter there would be no token until the switch was moved). However, in either case it is these tokens that form the basis of the smart home reasoning process.

Our interest in the smart home begins at this question of how to represent events and sensory signals by tokens in the most useful way. We are not directly interested in what the sensors are, or how they are processed, but rather how to extract useful data from this token stream and interpret that data; as is indicated in Figure 1 we see these as data mining and artificial intelligence problems, and all of the methods that we present here are from these areas. The more we have looked into the data that is presented in smart home problems, the more important context awareness and spatio-temporal reasoning have seemed to us. Actions and events do not occur in isolation, and things that seem odd as individual actions take on different interpretations once fleshed out with further information: a person balancing on one leg with their hand clutching the other leg is not unusual in the course of a yoga exercise program, but is otherwise rather unexpected.

Before we begin to describe methods that deal with these problems of spatio-temporal reasoning and context awareness, we describe our motivation for studying smart homes at all, which is centred around our target application area.

2 Motivation

A question often raised in connection with ambient intelligence, smart homes, context awareness, and the detection of human behaviour is whether the sacrifice of privacy is warranted by the benefits that the technology has to offer. The answer to this question lies in a problem that developed countries around the world is facing,

viz. the problem of an aging population. Life expectancy is higher than ever before, and so is the expectation of a high-quality, independent lifestyle throughout ones entire existence, irrespective of age or illnesses such as Parkinson's or Alzheimer's disease. Unfortunately, this expectation is not always met. Rather than moving an individual to a nursing home or employing the continuous support of a caregiver, an ambient intelligence can help people to live independently for longer by monitoring their behaviour in a non-obtrusive way, and alerting a caregiver if a critical situation arises.

We can separate the group of elderly people who desire to live an independent life into three groups. First, there are those with mobility and other problems associated with some sort of physical handicap (e.g., loss of a limb, severe arthritis). Second is a group who are of normal mental and physical health for their age. However, even healthy aging tends to be accompanied by at least some cognitive impairment (mainly involving memory problems and slowed processing speed). Third, there are people that suffer more than normal cognitive impairment for their age, usually involving dementia [29] (poor intellectual functioning involving impairments in memory, reasoning, and judgement).

Overall, between 40% and 60% of the independently living elderly suffer from some degree of cognitive impairment [37]. There exists a very large set of everyday domestic activities that this group may have difficulty carrying out. These include turning electrical or gas-driven appliances on and off, turning taps on and off, making tea or coffee, ironing, laundry, remembering to wash hands, and turning lights on and off. Quinn et al. found that bathing, personal hygiene, and getting dressed were major problems for their sample of 80 elderly people (average age 79). 82% needed assistance with managing their medication.

Many of the problems listed above can be monitored, and corrected, with a (relatively) unobtrusive ambient intelligence, and such a system could also alert caregivers to dangerous situations that it could not handle itself. Furthermore, in cases where two elderly people live together, one is often the sole caregiver for a less fortunate partner. This can be an incredibly stressful role for an aging caregiver, often resulting in serious illness, or even death [20]. An ambient intelligence, with its ability to monitor and intervene in many situations, could be of great assistance to an overworked, elderly caregiver.

The upshot of this discussion is that we are aiming for an unobtrusive monitoring system that does not necessarily interact directly with the inhabitants of the smart home. This may well reduce the resolution of the sensors that can be used, making the problem of identifying behaviours more difficult. However, it should also make the acceptance of the system rather easier for the home inhabitants; particularly if it does little more than passively observe and alert a caregiver of a potential problem, as this can be considerably less intrusive than unwanted check-up visits from an uninformed carer.

Note also that we are considering cognitive impairment, such as is typically found in Alzheimer's and similar degenerative diseases. This means that behaviours can change over time and the system therefore needs to adapt to identify what is occurring in a dynamic learning environment, something that is a challenge for many

artificial intelligence methods. Just adapting to a new activity is not the solution, because this activity might be abnormal, even if it is repetitive. Falling down should not be interpreted as normal behaviour, even if it is happening frequently, whereas preparing tea instead of the usual coffee is most likely not something a caregiver has to worry about.

There are two other potential problems with this type of smart home environment, which are rather common to all medical and health-related analyses. These are the vexed issue of false positives, and the fact that poor health outcomes are (fortunately) rather rare in the general population. False positives are a major reason why genuine emergencies are missed. In a modern day version of the boy who cried wolf, a machine that sets off an alarm every 10 minutes is ignored when the alarm is real. A system that alerts the caregiver twice a day to false positives is worse than having no system at all.

The other side of this coin is that the inhabitant of the house performing some unusual or potential dangerous behaviour is presumably rather rare, especially since the system is observing their activity 24 hours a day. In order to get a 99.99% success rate all that the smart home has to do is declare that the inhabitant never doing anything abnormal. From the application point of view, this obviously removes the need for the system at all. For machine learning methods based on real data it is a challenge to solve this problem, and there are two typical approaches. The first is to select the training data very carefully to ensure that only things of interest are seen, so that the system only learns to recognise things of real interest. There are several problems with this, including the huge amount of human involvement, and the fact that the system will not identify a potentially dangerous activity that was not introduced in training. In addition, the bias towards 'healthy' observations has another potential problem, which is that almost all of the training data shows people behaving normally, and therefore identifying abnormal behaviour, which is a potential indicator of ill health, is rather difficult.

We prefer to consider the intelligence system to be always learning, identifying for itself from the datastream when behaviours are potentially risky. The second approach, which is rather more useful to us, is to modify the error criterion so that success for the system is not simply the number of correct predictions that it makes. If the concept of risk is introduced, where the risk of annoying the caregiver can be balanced against the risk of missing a potential illness, then the success of the system can be more usefully measured, making it more effective. This can be done by something as simple as a 'loss matrix', which specifies the cost of making an incorrect prediction for the various possible outputs [7, 26].

Having highlighted these issues, we next describe a number of different techniques from the literature that may be useful in solving some of these problems, and discuss possible applications of them in the smart environment arena, beginning with methods to reason about human behaviours.

3 Sensory Input

While we are not directly interested in the sensory input, it is worth considering to what extent they provide direct input concerning spatial and temporal context. Based on our thesis that a smart home should be minimally obtrusive, we assume that the sensors will be minimally informative. Many of them could be as simple as location sensors, the use of electrical devices, etc. These sensors have the benefit of being cheap and easy to install, but they usually provide very limited information compared to, for example, a camera. This has the benefit that very limited preprocessing is required, but the disadvantage that not much information is available for the ambient intelligence. Consider the difference between an image of a person moving round the kitchen, filling and boiling the kettle, placing a teabag in a cup, adding milk, and pouring boiling water on, as opposed to the fact that electricity was used from a particular plug point and the fridge opened.

The next question is precisely what form the data should be presented in. We assume that the system has some information about the location of the sensor (at least, the room it is located in) and the time that it was activated at, but there are a variety of different options for how the sensor data can be presented. At one extreme of the options is that the intelligence periodically polls each sensor and gets a reading from it, so that the sensors are always reporting their state, even when there is nobody in that particular room. In this case, the output from a power point could be a series of readings of the values of the power drawn from the socket. While this could be useful (since it might enable the actual device being used to be identified), we prefer to think that this kind of processing is built into the sensor, and the output that is presented is a token that represents the fact that a device (whether specified or not) is being powered from that particular socket.

Given that we are assuming that tokens are presented by each sensor, there are still different options, but our default hypothesis is that each sensor presents two different tokens, corresponding to ‘sensor has started’, ‘sensor has stopped’ that are presented in an interrupt-like fashion by the various sensors. While this may not be suitable for certain types of sensor, we generally assume that it is useful.

4 Spatio-Temporal Reasoning

As was pointed out before, interpreting human behaviour in context involves reasoning about space and time. For example, preparing a meal at noon in the kitchen is usually a perfectly normal behaviour, but if the same activity occurs at 3 in the morning in the garage, then it is a behaviour that needs some special attention.

When referring to spatial and temporal reasoning, we mean reasoning in a similar way to that in which a person would deal with time and space in everyday life, and not the way how a physicist, for example, would describe locations, regions, time points, intervals, events etc. and reason about them. It has been argued in AI that the layperson’s everyday form of spatio-temporal reasoning is of a qualitative,

rather than a quantitative nature. We are not usually interested in precise descriptions of space and time. Coarse and vague qualitative representations frequently suffice to deal with the problems that we want to solve. For example, to know that the meal is prepared in the kitchen (rather than knowing the exact coordinates for this activity) and at lunchtime (rather than 12:03:37) is often enough to decide whether the behaviour is normal or not.

Most qualitative approaches to spatial and temporal reasoning are based on relations between objects (such as regions or time intervals). They are fundamentally different from the approaches developed in the area of computational geometry [5, 36]. Of the temporal qualitative approaches, two are dominant and reappear frequently in different scenarios: Allen's temporal logic [2] and the point algebra [50]. The first one uses intervals to describe specific events, and relations between intervals to describe interdependencies between events (see Figure 2). For example, if the event of the door to the room being opened occurs before the event of the person entering the room, and if I_1 denotes the time intervals in which the door-opening event occurs and I_2 the one for the person-in-room event, then we would have the relation $I_1 < I_2$.

Together with the 13 relations depicted in Figure 2, Allen introduced an algorithm that uses a composition table to reason about networks of relations. The table determines the possible relations between two intervals like I_1 and I_3 given the relations between I_1 and another interval I_2 as well as the relation between I_2 and I_3 (see Figure 3). If, for example, the door to the room is opened (I_1) before the person is in the room (I_2) and the TV is on (I_3) during the person's presence in the room, then the composition table tells us that the door is opened before the TV is switched on:

$$I_1 < I_2 \text{ and } I_2 \text{ di } I_3 \text{ then } I_1 < I_3.$$

If the TV being on overlaps with the person being in the room (it might have been on before the person entered the room, but is switched off while the person is in the room), then the door to the room being opened either overlaps with the TV being on, is finished by the TV being on, or contains the event of the TV being on:

$$I_2 \text{ di } I_3 \text{ and } I_3 \text{ o } I_3 \text{ then } I_1 \{ \text{o, fi, di} \} I_3.$$

A set of possible relations such as {o, fi, di.} is also called a non-atomic Allen relation. They provide a way to express uncertainty. In this case, we only know that one of the relations holds, but we are not able to tell which one it is, given the current information about the events. When at a later stage more information becomes available, we might be able to resolve the uncertainty.

Unlike Allen's temporal logic, the point algebra uses time points and three possible relations to describe interdependencies among them (see Figure 4). These relations can be used in a similar way as interval relations. For example, the time point at which we start to open the door precedes the time point at which we commence entering the room. In fact, Vilain and Kautz pointed out that many interval relations can be expressed as point relations by using the starting and finishing endpoints of

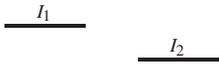
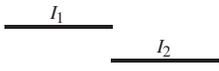
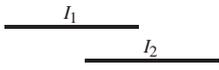
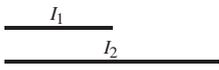
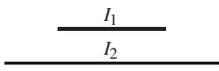
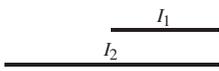
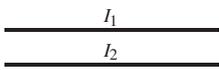
<i>Relation</i>	<i>Illustration</i>	<i>Interpretation</i>
$I_1 < I_2$ $I_2 > I_1$		I_1 before I_2 I_2 after I_1
$I_1 m I_2$ $I_2 mi I_1$		I_1 meets I_2 I_2 met by I_1
$I_1 o I_2$ $I_2 oi I_1$		I_1 overlaps I_2 I_2 overlapped by I_1
$I_1 s I_2$ $I_2 si I_1$		I_1 starts I_2 I_2 started by I_1
$I_1 d I_2$ $I_2 di I_1$		I_1 during I_2 I_2 contains I_1
$I_1 f I_2$ $I_2 fi I_1$		I_1 finishes I_2 I_2 finished by I_1
$I_1 = I_2$		I_1 equals I_2

Fig. 2 Allen’s thirteen atomic relations.

the intervals. The interval relations that can be expressed in terms of point relations are called pointisable interval relations.

Both Allen’s interval logic and the Villain and Kautz’s point algebra can be used for spatial reasoning by interpreting intervals as one-dimensional objects and time points as locations in space, respectively. For example, if a person (O_1) is standing left of a box (O_2) and if the box contains a book (O_3), then we can conclude that the person is standing left of the book:

	<	m	o	fi	di	si	=	s	d	f	oi	mi	>
<	<	<	<	<	<	<	<	<	<	<	<	<	<
m	<	<	<	<	<	m	m	m	o, s d	o, s d	o, s d	f, = fi	di si, oi mi, >
o	<	<	<, m o	<, m o	<, m o fi, di	o fi, di	o	o	o, s d	o, s d	o, s d, f, = fi, di si, oi	di si, oi	di si, oi mi, >
fi	<	m	o	fi	di	di	fi	o	o, s d	f, = fi	di si, oi	di si, oi	di si, oi mi, >
di	<, m o fi, di	o fi, di	o fi, di	di	di	di	di	o fi, di	o, s d, f, = fi, di si, oi	di si, oi	di si, oi	di si, oi	di si, oi mi, >
si	<, m o fi, di	o fi, di	o fi, di	di	di	si	si	s, = si	d, f oi	oi	oi	mi	>
=	<	m	o	fi	di	si	=	s	d	f	oi	mi	>
s	<	<	<, m o	<, m o	<, m o fi, di	s, = si	s	s	d	d	d, f oi	mi	>
d	<	<	<, m o, s d	<, m o, s d	<, m o, s, d f, =, fi di, si, oi mi, >	d, f oi, mi >	d	d	d	d	d, f oi, mi >	>	>
f	<	m	o, s d	f, = fi	di si, oi mi, >	oi mi, >	f	d	d	f	oi mi, >	>	>
oi	<, m o fi, di	o fi, di	o, s d, f, = fi, di si, oi	di si, oi	di si, oi mi, >	oi mi, >	oi	d, f oi	d, f oi	oi	oi mi, >	>	>
mi	<, m o fi, di	s, = si	d, f oi	mi	>	>	mi	d, f oi	d, f oi	mi	>	>	>
>	<, m o, s, d f, =, fi di, si, oi mi, >	d, f oi, mi >	d, f oi, mi >	>	>	>	>	d, f oi, mi >	d, f oi, mi >	>	>	>	>

Fig. 3 Allen’s composition table. The entry at row r_1 and column r_2 in the table denotes the possible relations between O_1 and O_3 , assuming that $O_1r_1O_2$ and $O_2r_2O_3$.

$$O_1 < O_2 \text{ and } O_2 di O_3 \text{ then } O_1 < O_3.$$

In this example, we treat a horizontal axis that serves as a reference frame the same way we treat the time axis.

For Allen’s interval logic, several multi-dimensional extensions have been suggested, which use tuples of intervals that represent projections of regions onto the axes of a Cartesian reference frame [14, 32, 4]. These representations capture the topological and orientational aspects of space. For example, consider the sketch in

<i>Relation</i>	<i>Illustration</i>	<i>Interpretation</i>
$P_1 < P_2$	P_1 P_2 ● ●	P_1 precedes P_2
$P_1 = P_2$	P_1 ● P_2 ●	P_1 same as P_2
$P_1 > P_2$	P_2 P_1 ● ●	P_1 follows P_2

Fig. 4 The three different point relations from the point algebra.

Figure 5, which shows three boxes arranged on a flat surface such as a carpet. When

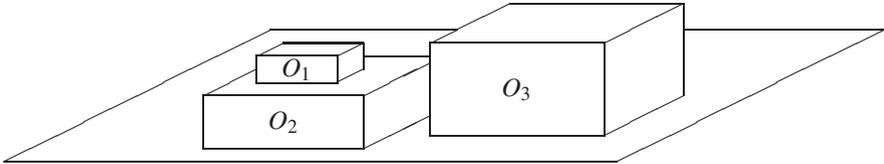


Fig. 5 Three boxes arranged on a flat surface.

describing the spatial relations between the boxes, we can split each relation into three components, each of which describes an aspect of the relation with respect to a particular axis: the first component with respect to a horizontal axis, the second with respect to a vertical axis, and the third with respect to a depth axis:

$$O_1 \langle d, mi, d \rangle O_2 \text{ and } O_2 \langle <, s, o \rangle O_3.$$

Applying the transitivity table to the components of the relation triples yields an additional spatial description:

$$O_1 \langle <, \{d, f, oi\}, \{<, m, o, s, d\} \rangle O_3.$$

We argue that the topological aspects of space are the most relevant ones for modelling human behaviour in an ambient intelligence. For example, to determine what activity a person might be involved in, it is useful to know which room that

person is in (although not necessarily the person's exact position in that room). Or, to determine the possible movements of a person, it is useful to know which rooms are connected to which other rooms. We will therefore focus on the topological aspects of space in the following, introducing one topological spatial reasoning calculus as an example: the Region Connection Calculus or RCC [39].

The basis of the RCC theory is a reflexive and symmetric relation, called the connection relation, which satisfies the following axioms:

1. For each region X : $C(X, X)$
2. For each pair of regions X, Y : $C(X, Y) \rightarrow C(Y, X)$

From this relation, additional relations can be derived, which include the eight jointly exhaustive and pairwise disjoint RCC8 relations: listed below and illustrated in Figure 6.

$$\text{RCC8} = \{\text{DC}, \text{EC}, \text{PO}, \text{EQ}, \text{TPP}, \text{TPPi}, \text{NTPP}, \text{NTPPi}\}.$$

Reasoning about space is achieved in the RCC theory by applying a composition table to pairs of relations, similar to the composition table in Allen's interval logic. Given the relation R_1 between the regions X and Y , and the relation R_2 between the regions Y and Z , the composition table determines the relation R_3 between the regions X and Z . For example, if the region that the person is standing in (R_1) is disconnected from the region occupied by the box (R_2) and if the region taken up by the book (R_3) is a tangential proper part of R_2 , then we can conclude that R_1 and R_3 are disconnected:

$$\text{DC}(R_1, R_2) \text{ and } \text{TPPi}(R_2, R_3) \text{ then } \text{DC}(R_1, R_3).$$

The RCC theory is not the only approach to reasoning about topological relations. Egenhofer and Franzosa [10] consider the intersections of boundaries and interiors of pairs of regions and derive a formalisation from that, called the 4-intersection calculus (and by adding the complement of the regions to the intersections, the 9-intersection calculus).

5 Symbolic Approaches

Since its beginning, the area of artificial intelligence has been split into two camps: symbolic and subsymbolic. Researchers from the symbolic camp base their work on the physical symbol system hypothesis, which states that a "physical symbol system has the necessary and sufficient means for intelligent action" [33]. In the context of detecting behaviour from patterns, this means the explicit representation of actions and events, together with the ramifications of these. Although there are various ways to achieve this, one approach has gained frequent attention in AI: the situation calculus [30]. This calculus is a logic-based approach that uses a state description in its formulas to describe what happens when a particular action is

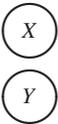
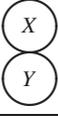
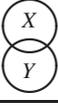
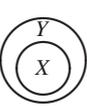
<i>Relation</i>	<i>Illustration</i>	<i>Interpretation</i>
$DC(X,Y)$		X disconnected from Y
$EC(X,Y)$		X externally connected to Y
$PO(X,Y)$		X partially overlaps Y
$EQ(X,Y)$		X identical with Y
$TPP(X,Y)$		X tangential proper part of Y
$TPPi(X,Y)$		Y tangential proper part of X
$NTPP(X,Y)$		X nontangential proper part of Y
$NTPPi(X,Y)$		Y nontangential proper part of X

Fig. 6 An illustration of the RCC8 relations.

performed. For each action, an axiom is provided that specifies when the action is possible, and another axiom that specifies the effects of the action. For example, if the kettle is not empty in a situation s , then it is possible to heat the water in the kettle. The heating action results in a new situation in which the kettle contains hot water. Formally:

Possibility axiom: $Kettle(K) \wedge \neg Empty(K, s) \Rightarrow Poss(Heat(K, s))$
 Effect axiom: $Poss(Heat(K, s)) \Rightarrow Hot(K, Result(Heat(K, s)))$.

The function *Result* denotes the situation that results from situation *s* if the kettle is heated.

Functions and predicates that can change from one state to another are called fluents (e.g., the predicate *Hot*), whereas predicates and functions that don't change are called atemporal or eternal (e.g., *Kettle(K)*). The effect axioms specify which fluents change from one state to another when a particular action is performed. They do *not* specify the fluents that stay the same. The problem of representing these is commonly referred to as the frame problem, and it can be addressed by writing frame axioms:

Frame axiom: $Door(D) \wedge \neg Open(D, s) \Rightarrow \neg Open(D, Result(Heat(K, s)))$.

The situation calculus is commonly used to perform planning tasks. Given a set of possible actions, a constructive inference engine can then be used to find a sequence of actions that achieves a desired effect. For example, given actions like filling a kettle with water, heating the kettle, putting a tea bag into a cup, pouring water into the cup, etc. the system can assemble them in such a way that it results in achieving the task of making tea. For the purpose of this chapter, we are interested in the inverse problem, which is referred to as the projection task: given a sequence of actions, what is the outcome of the sequence?

Although, at first glance, the projection task seems to be trivial, unfortunately this is not the case if we attempt to go beyond purely describing the state of the fluents. For example, it is not very useful to know that after performing the aforementioned actions, the kettle is empty, the tea bag is in the cup, the cup is filled with water, etc. What we rather want is associated a symbol, in this case *TeaMaking*, with these actions.

A straightforward solution to this problem is keeping a database that associates action sequences with their interpretations. Such an approach is feasible if each interpretation is caused by one or very few sequences of actions, which generally is not the case. There are many ways to make tea, like using a pot on the stove rather than a kettle, using loose tea leaves instead of a tea bag, replacing part of the water with rum, etc. Moreover, a sequence of actions can be interleaved with another sequence of actions. For example, while waiting for the water to boil, we might want to prepare a sandwich to go with the tea.

Another shortcoming of this approach is the lack of an explicit time and space representation. Although it is possible to use temporal or spatial references instead of states in the axioms of the calculus (like $Happens(Heat(K), [10:00, 10:10])$) and thereby turning the situation calculus into an event calculus, the resulting calculus still does not automatically inherit the reasoning techniques known for Allen's temporal algebra or the Region Connection Calculus. Bhatt and Loke [6] have suggested a different approach, which keeps the causal theory (represented in the situation calculus) separate from the spatial theory (represented, for example, in RCC-8). This approach provides a cleaner and more powerful model, but it also puts an extra burden onto the modeller who has to develop the theory.

An approach based on the situation calculus assumes that there is a reasoning engine that interprets the detected sequence of actions and makes an assessment of whether it constitutes a normal or abnormal behaviour. Augusto and Nugent [3] suggest a different approach, which explicitly models abnormal behaviour in the form of Event-Condition-Action rules:

ON event IF condition THEN action.

The system that uses these rules keeps a current state, which is updated whenever an event occurs. If the condition associated with the event is satisfied, a particular action is triggered, which can be raising an alarm, notifying a caregiver, or just noting a change in the state of the system.

The rule-based approach provides explicit control over the detection of abnormal behaviour and the action to be taken if such a behaviour occurs. However, it still requires the modeller to create a reasonable set of rules that accounts for all forms of abnormal behaviour. The approaches discussed in the next section offer a solution to this problem by having the system learn to distinguish between normal and abnormal behaviour.

6 Machine Learning Approaches

In addition to the symbolic approaches identified previously, there are also significant contributions that subsymbolic and probabilistic methods can make to the area of context awareness in smart homes; subsymbolic and probabilistic methods are often combined under the generic name of ‘machine learning’. In this chapter, three machine learning methods that can be used for spatio-temporal learning and context awareness in the area of ambient intelligence are described. The three methods provide potential solutions to different parts of the smart home problem, as will be identified as they are described. The section finishes with a brief discussion of other machine learning methods that might be of interest for spatio-temporal learning in smart homes.

6.1 Frequent Pattern Data Mining

Assuming that the sensory signals have been encoded in some way into a set of tokens that represent certain events (or the cessation of events), as was discussed previously, the first question is how to identify patterns in this sequence. The token stream is temporal in the sense that tokens appear when events trigger them, but they can be assembled into a sequence and this can be mined as a non-temporal sequence.

For example, the token signal over each hour, or some other time window, could be assembled into a string, and then these strings can be treated as separate ‘transac-

tions'. The window could be strictly temporal in size, or of a fixed number of tokens (to see the difference, consider that a time window of one hour will probably have rather more tokens in it at 6 pm than at 3 am). Selecting a suitable size of window is important, and not trivial. It can be done based on the data using a technique such as cross-validation [7, 26]. The problem for the ambient intelligence is to identify from these strings those patterns of tokens that represent particular behaviours. One way to consider the problem is to say that tokens that appear together frequently are likely to be parts of the same behaviour, or parts of related behaviours. We can therefore appeal to a variety of data mining procedures that aim to do so-called market basket analysis. The aim of these methods is to extract combinations of items that occur together frequently. A typical motivation for these algorithms is to identify patterns of shopping, such as that pasta and red wine are often bought at the same time, or CDs by certain artists, so that shop layout and advertising can be better targeted to customers' needs.

One method of performing frequent pattern mining is the FP-tree [15], which is a tree data structure that extracts patterns with support (i.e., appearance in the dataset) above some pre-defined threshold. It was developed as an alternative to the Apriori algorithm of Agrawal and Srikant [1], and consists of a tree with individual items at each node, with the nodes at the top of the tree having the highest support. There is also a header table that indexes individual items so that they can be quickly found in the tree (usually implemented as a linked list if an item appears more than once within the tree). As an example, suppose that three different sets of tokens are seen in three windows: $\{A, C, D\}$, $\{A, B, D, E\}$, $\{B, D, E, F\}$. The algorithm first counts the support of each token individually, and sorts them into order of decreasing frequency, which would be $D : 3, A : 2, B : 2, E : 2, C : 1, F : 1$. Assuming that the minimum support (a user-defined parameter) was chosen as 2, those items with support 1 would be removed (here C and F) and then the first record ($\{A, C, D\}$) would be put into the tree, starting with D since it has the highest support. C would be ignored since it has support below the minimum, and therefore any pattern that it is part of would have support below the minimum. Each node is annotated with its support. The tree after the first pattern has been entered is shown on the left of Figure 7, and the tree after all three patterns have been entered is shown on the right of the figure.

Many methods of data mining can be understood in terms of compression and related information-theoretic ideas. If the cost of transmitting or storing data is to be kept to a minimum, then a pattern that occurs frequently should be represented by a short bit sequence in a codebook, while one that occurs less commonly can have a longer one without incurring much additional cost, since it is seen less often. In this way, the problem of identifying frequent patterns can be mapped into one of finding the shortest bit sequence that encodes the information. The means that statistical measures of complexity such as the Minimum Description Length [41] are of interest, as in [13, 16, 9].

There will be two aims in mining frequent patterns: to identify the longest repeated sequences, and to identify the sequences that occur the most often. For be-

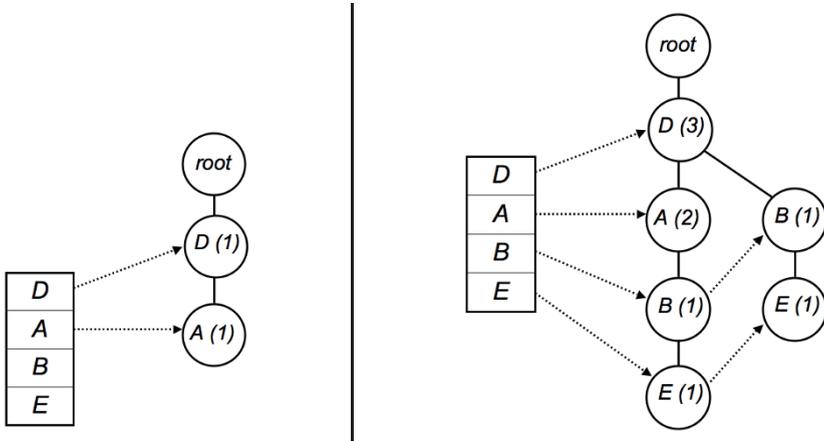


Fig. 7 *Left:* The FP-tree after the first pattern is entered. *Right:* The FP-tree after all three patterns have been entered.

behaviour identification we are most interested in the first of these, although the second can also provide useful information.

One thing that makes the task of behaviour identification significantly more difficult is that there is no reason why patterns in a behaviour should be contiguous in the sequence; for example, when making tea, while waiting for the kettle to boil a person can perform other unrelated tasks such as reading the newspaper or taking out the rubbish. This problem is radically more difficult if there is more than one person in the house. Additionally, there is no reason to assume that tasks are always performed in the same order. Again, tea making provides an example; whether you pour the milk in first or last is a matter of preference, choosing one or the other does not mean that you are performing a different qualitative behaviour. Of course, some people take no milk at all, and other add sugar, and so the actual tokens that comprise one description of the behaviour can vary slightly, as well as their order.

A possible way to deal with the first problem could be to completely remove the temporal sequence by sorting the tokens into some other ordering (such as alphabetical based on the arbitrary naming of the tokens) and to use that. This is what is commonly done in data mining, but it has two disadvantages here: firstly that a token that appears in several behaviours and is thus seen several times in the string will probably only be shown once, and the remainder suppressed, and secondly that the temporal information can be useful to discriminate between different behaviours. For example, if a person were to boil the kettle, take milk from the fridge, and remove a cup from the cupboard, then they could be making a cup of tea, but if the first action happened 35 minutes after the other two, that would be rather unlikely; rather they are likely to be events in unrelated behaviours. A possible solution to this problem is presented by Tavenard et al [48], based on the T-pattern of Magnusson [24]. This uses the time interval between events to try and recognise when particular events are related and when they are not.

This is not to say that frequent pattern data mining is not useful here; it may well be that by identifying groups of tokens that represent particular behaviours, and presenting them in such a way that spatio-temporal reasoning can be applied to the output is a very useful form of preprocessing. By removing the various temporal aspects of the data (except that some kind of windowing is applied to the data stream, so that only tokens that appear relatively close together in time could form part of a pattern) it is possible to identify actions that frequently happen at similar times and that could therefore be part of the same behaviour, without having to consider the temporal order that they appear in.

There are multiple challenges in this area, some of which have been identified above, and it is an interesting area of study. It may prove to be a fruitful approach to the problem of identifying behaviours at a relatively naïve level. The question then is how to use the identified behaviours, and the tokens that comprise them, in order to recognise whether or not a person is behaving normally. Several methods for doing this have already been described, but in the next two sections we will now identify two techniques that may well have their part to play in dealing with this challenge.

6.2 Graphical Models

One of the most popular areas of research in machine learning currently is that of graphical models. These are a general framework where variables are identified as nodes and edges between the nodes represent conditional relationships between nodes (more correctly, the absence of an edges implies conditional independence of the two nodes). Graphical models have been described as a marriage between computational graph theory and probability, and many common machine learning algorithms can be represented within this framework, including the Kalman filter [19], Hidden Markov Model [38], and Bayesian Network [18]. The graph that underlies the model can be directed (i.e., the edges have arrows on) or undirected. In this second case the models are known as Markov random fields and have shown particular application to image analysis.

In this chapter we will restrict ourselves to discussing one particular example, which is to take sequences of tokens that have been identified as forming a behaviour (possibly using the data mining methods described above), and to use examples of when these behaviours occurred in order to attempt to identify these behaviours in the stream of tokens as they are produced. In this way, behaviours can be identified in real-time, and significant departures from normal behaviour hopefully recognised before they become dangerous. Since we are interested in these sequences as they appear in time, we are looking at dynamic Bayesian Networks, a subclass of graphical models that includes Hidden Markov Models (HMM) as its most useful and computationally tractable example. Hidden Markov Models have met with considerable success in another area of temporal sequence analysis, that of speech recognition [38].

The important point about HMMs is the disconnection between observations, which for our smart home application are the tokens that are provided by the sensors, and the underlying states that gave rise to the observations. For example, observations could be that water is coming out of the tap, and possible states that give rise to that are that somebody is using the tap, that it leaks, or that there has been a sudden drip after use.

As an example of a possible sequence of observations and the states that could be represented by them, Figure 8 shows a possible sequence of tokens and the states that could have given rise to them (in the figure, token *F* is about the tap being used, *C* and *D* are concerned with electricity being used, and *X* and *Y* with the fridge being opened and closed). The figure does not show the probabilities.

Linking the sequence of observed tokens over time into lists of actions that can eventually be recognised as behaviours is still non-trivial, however. For example, from a stream of observations, there may be some that are not connected to the current activity, and the order in which the events occur can vary. Normal Hidden Markov Models will not deal well with either of these cases, and extending the paradigm to deal with these and other aspects of temporal sequences is an area of current research. One approach that shows some promise for dealing with the ordering problem is the Conditional Random Field [46], which encodes conditional probabilities of sequence order in an undirected graphical model. While these methods show some promise in removing some of the application-oriented limitations of the Hidden Markov Model, they do incur costs, such as the fact that Markov Chain Monte Carlo (MCMC) methods are needed to approximate the probability distributions, rather than the expectation-maximisation (EM) algorithm of the Hidden Markov Model [26].

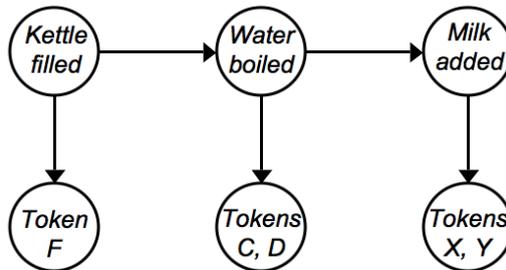


Fig. 8 A sample Hidden Markov Model for recognising tea making behaviour. Note that it is not sufficient since it does not include the possibility that the milk is added before the water is boiled.

One perceived limitation of the HMM is that each state only produces one observation. While this makes sense for many applications, it is not intuitive for behaviour recognition. A suggestion of one solution is provided in Figure 8, where two tokens together are identified as a single observation. However, there are other solutions possible, such as modifying the structure of the HMM to allow for more complex

models. The Hierarchical HMM [12], where each state is allowed to be a complete hierarchical HMM of its own is one such variation that allows complete activities to be built into each of the states. This has been applied to human behaviour recognition by Nguyen et al [34].

Our own focus in this area is to look for a simple solution. By considering the problem of identifying and predicting behaviours as one of competition between a set of Hidden Markov Models, each of which recognises one particular behaviour, the best match for the current set of tokens can be identified probabilistically. In the event that none of the models produces a good match, either the system needs to add a new model and train it on this data, or it needs to alert the caregiver.

6.3 Novelty Detection and Habituation

A well-publicised criticism of neural networks and other machine learning methods is that they will always give you an answer, even when the input that is presented is significantly different from those that they were trained on. There are a variety of ways to solve this problem, but one that is of particular interest is novelty detection. In a novelty detection system the algorithm is trained on ‘normal’ behaviour that is seen during the training phase, and then in use, the algorithm classifies signals as either normal or novel.

The application to smart homes is most clear in the monitoring scenario that we are considering. Assuming that methods have been used to process sensor signals into a pattern of human behaviour, the task is to decide whether or not the person is acting normally. Some early work in this area is reported by Rivera-Illingworth et al [42] and Jakkula and Cook [17]. In monitoring of life activities, abnormal behaviour could be a completely new activity (jumping up and down aimlessly), performing a normal behaviour in an unexpected place (watering the TV, not the houseplants) or unexpected time (drinking vodka in the morning), or an unexpected person performing a behaviour that, while common, they have not done before (the staid aunt who removes her clothes and wanders around naked). As O’Keefe and Nadel [35] phrase it (page 241):

“...the novelty of the wife in the best friend’s bed lies neither in the wife, nor the friend, nor the bed, but in the unfamiliar conjunction of the three.”

Of the four types of novelty identified, only the first lacks context as part of what makes it novel, and this highlights clearly how important context (whether in the form of the person, the place, or the time) is in identifying and categorising behaviours.

There are many methods of performing novelty detection described in the machine learning literature, from the multi-layer Perceptron-based novelty filter of Kohonen and Oja [22] to methods based on Support Vector Machines [8]; for an overview of methods, see [25]. However, here, rather than looking at a method based

on a particular machine learning algorithm, we will see how a simple model of the biological phenomenon of habituation can be used to perform novelty detection.

Habituation is widely recognised as one of the simplest forms of learning in biological organisms. It consists of a reduction in response rate to a stimulus that is presented repeatedly without ill effect. This has important survival benefits for animals, as it enables them to ignore unhelpful but common stimuli in order to focus on potentially more important things. Habituation has been identified in organisms as simple as the sea slug *Aplysia* and the nematode worm, and as complex as humans; we habituate to stimuli constantly, from the way that we stop noticing the sensation of the socks on our feet to the ignoring of background sounds, such as those produced by machinery. The effects of habituation are often apparent in psychology experiments, and care must be taken to avoid confounding results.

Thompson and Spencer [49] produced a list of nine identifying characteristics of habituation; these were recently revised by an interdisciplinary group of researchers convened for the purpose, and their characteristics, and the justification for the changes from Thompson and Spencer [49] are described in Rankin et al [40].

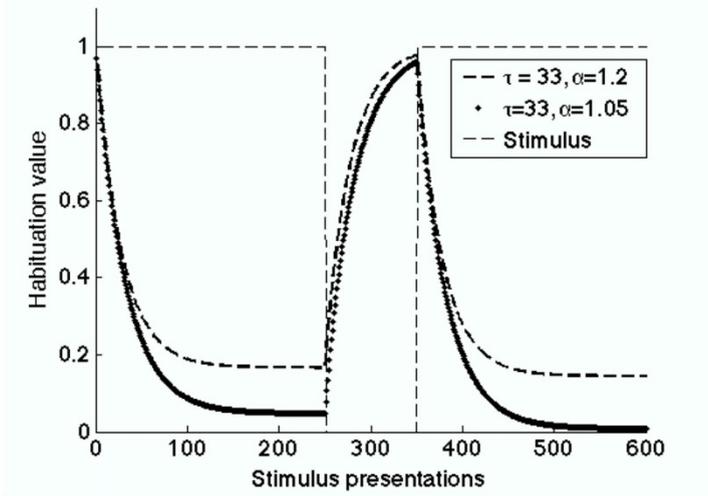


Fig. 9 Habituation causes an exponential decay in response with stimulus presentation that recovers when the stimulus is withdrawn.

The effect of habituation on the response to a particular stimulus is generally seen as an exponential decay of the response strength that recovers when the stimulus is not presented, as shown in Figure 9. In addition, there can be spontaneous dishabituation when a different stimulus is presented. Modelling this computationally can be as simple as a negative exponential function [44], although models are typically more complex than that. A recent review of modelling habituation and its application in machine learning is provided by Marsland [27].

For our smart home application area, to demonstrate the use of habituation, we assume that behaviours are encoded in some way, possibly as a set of identified features that categorise individual behaviours. Temporal and spatial information could be added as additional input features. These can then be classified using any clustering algorithm, such as the Self-Organising Map [21] or GWR network [28]. The outputs of the network can be modified by using the habituation model, so that inputs that are seen frequently are habituated to, while those that are seen only rarely are not. The network is then trained on data collected from normal living in the smart home; simply the typical data of daily life, either specific to a particular individual, or collected over many individuals in different houses.

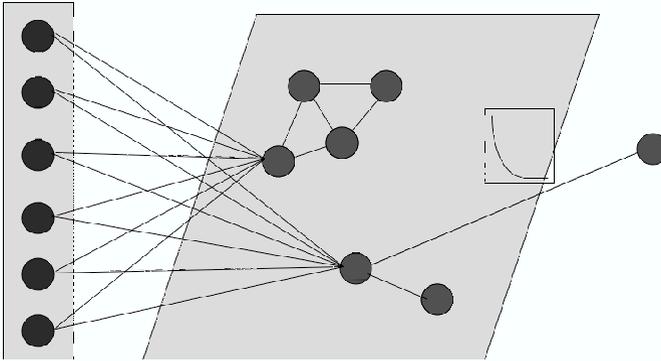


Fig. 10 Combining a model of habituation with a neural network requires modifying the output of the neural network by the amount of habituation that has already occurred for that input class. When a class is seen often its class is habituated to, whereas for a rare class the output is strong.

After training, the output of the habituation will suggest whether a particular input behaviour is normal or has novel features, either in the behaviour itself or its spatial or temporal features, in relation to the training set. This model can be modified to provide additional information, such as distinguishing between behaviour that is normal because it has been seen in the past, but not recently, and completely typical behaviour. The way in which this can be done is to use two different habituating responses. One of them learns slowly, but does not forget; when an input is seen the response decrements for each presentation, but not much. However, when the input is not seen the strength does not change. The other one learns more quickly, so that fewer presentations are required for an input to be normal, but it forgets, so that if an input class is not seen, the habituated response recovers to its maximum level. Consider then the effect of an input that is seen frequently. It will have both outputs habituated and will therefore produce no response. However, an input that has been seen in the past, but not recently, will produce output from one of these habituation modules, but not the other. Interestingly, this setup also enables a second interesting

case to be identified, which is a significant and sustained change in behaviour; when a new input is seen a few times the rapidly learning module will habituate to it, but the long-term one will not. Thus, an input that produces an output only from the long-term module is something that was rare, but is now being seen frequently.

It was suggested above that spatial and temporal context can be encoded as inputs to the learning algorithm. However, there are alternatives to this choice, such as training banks of learners, each of which are focused on specific combinations of times and places. In this way, the context of each behaviour is explicitly defined only within that context.

6.4 Other Methods

Much of the learning that occurs in the smart home environment does not have clear targets. For example, the activities that the home should recognise are not necessarily described by humans, and we want to avoid having humans describe these in advance. This means that supervised learning methods are not appropriate for many of the tasks. The alternative that is used in the novelty detection methods described above are to use no external feedback and only expect the neural network to cluster relevant data together and detect outliers from it. However, there is an inbetween approach known as reinforcement learning [47], where an external signal is provided that gives feedback on whether or not the output was correct in the form of a reward, but not how to improve it. This makes the problem one of explicit search. One of the main challenges of reinforcement learning is known as the credit assignment problem, and consists of working out which actions led to the reward, since it may well not be those that were closely linked in time. For example, suppose that the smart home predicts that a person is making tea. The reward for this will be given once the observation of tea drinking is made, but the tokens that led to the prediction, such as running the tap and boiling the kettle occurred long before this observation was made. This problem has been considered in a simple smart home context for lighting control by Mozer [31].

Standard neural networks can also be adapted to perform spatio-temporal learning. It is the problem of time-series prediction that has most commonly been considered, although this can be dealt with by using inputs taken within a fixed-timewidth window, which means that the learning algorithm does not need any information about the time series. However, it is also possible to build temporal information into the learning algorithm explicitly. For example, the multi-layer Perceptron [43], can be modified to deal with temporal data in a variety of different ways: typically by using outputs from the network at previous timesteps as inputs (this is known as a recurrent neural network; there are a variety of algorithms, but [11] is a typical reference) or by modifying the behaviour of the neurons, so that they do not completely lose their activation between inputs. These are generically known as time delay neural networks, and an example is the leaky integrate-and-fire model of Stein [45]. For a review of the area of spatio-temporal neural networks, see Kremer [23].

7 Conclusion

The development of smart environments, such as smart homes, poses challenges to researchers from a wide range of areas, including engineers, who face the problem of finding adequate sensors to perceive the environment and effective actuators (including smart user interfaces) to act upon it; cognitive scientists, who seek to understand human behaviour and with that an understanding of the sensor data; and computer scientists, who endeavour to implement an ambient intelligence that can reason about the information perceived and can infer the appropriate actions from it. In this chapter we have only touched the tip of the iceberg by focusing on a very narrow aspect of the overall problem: the classification of complex human behaviour in a spatio-temporal context. In particular, we have focused on the problem of context awareness, where data concerning other activities taking place at approximately the same time, or in approximately the same place, require actions to be considered as a group, rather than in isolation. We have provided examples of this in the smart home context.

The chapter has most certainly fallen short of presenting one homogeneous solution to the problem of classifying complex human behaviour. This problem has been fundamental to the field of artificial intelligence since its beginning, and despite massive efforts for more than sixty years has not been solved in general. It therefore does not come as a surprise that we have not introduced a ‘universal ambient intelligence for smart environments,’ but rather have restricted ourselves by shedding some light on various aspects of the problem.

One aspect is to automatically find meaning in sensor data and to distinguish between normal and abnormal behaviour on the basis of what has happened in the environment in the past. The reason for this approach is obvious: without the need for explicit modelling of human behaviour, we can in principle more easily create a system that is customised to a particular environment and the people living in that environment. Such an approach should also be able to cope with the problem of a very large number of potential behaviours, as it would learn only those behaviours that are actually occurring in a particular setting: there are many different ways of preparing a cup of tea or making a meal, but each of us usually does so in only a very small number of different ways.

Despite having these appealing advantages, subsymbolic approaches have their shortfalls, one of which is the lack of associating meaning with sequences of events. It might be easy enough to detect a tea making sequence, but this does not mean that the system knows much about tea, such as that tea making is a way to provide a drink, that providing a drink is the result of a need for nourishment or satisfying thirst, that the satisfaction of such a need is essential to maintaining a healthy condition, and so on. It would therefore not be able to conclude that cutting up juicy fruit is similar to making tea, while washing the car is not (although on the surface it might seem the other way round, as washing the car involves turning on the tap while cutting up fruit does not). Explicit modelling on a symbolic level, involving ontologies and commonsense reasoning, can overcome this problem, given enough resources.

We hope to have provided a convincing argument that neither symbolic nor sub-symbolic approaches alone can provide a comprehensive solution to the problem of creating smart homes. However, we believe that many of the challenges in the development of smart homes can be overcome by carefully merging symbolic and sub-symbolic approaches and linking of symbolic reasoning methods with prediction and inference made through machine learning. It is in this direction that our own research is moving.

Acknowledgements

We are grateful for discussions and collaboration with the other members of the Massey University Smart Environments (MUSE) group (<http://muse.massey.ac.nz>).

References

- [1] Agrawal R, Srikant R (1995) Mining sequential patterns. In: Yu P, Chen A (eds) Proceedings of the Eleventh International Conference on Data Engineering, pp 3–14
- [2] Allen J (1983) Maintaining knowledge about temporal intervals. *Communications of the ACM* 26:832–843
- [3] Augusto J, Nugent C (2004) The use of temporal reasoning and management of complex events in smart homes. In: Proc. ECAI-04, Valencia, Spain, pp 778–782
- [4] Balbiani P, Condotta JF, del Cerro L (1998) A model for reasoning about bidimensional temporal relations. In: Proc. KR-98, Trento, Italy, pp 124–130
- [5] de Berg M, van Kreveld M, Overmars M, Schwarzkopf O (1997) *Computational Geometry*. Springer
- [6] Bhatt M, Loke S (2008) Modelling dynamic spatial systems in the situation calculus. *Spatial Cognition and Computation* 8(1&2):86–130
- [7] Bishop C (1995) *Neural Networks for Pattern Recognition*. Clarendon Press, Oxford
- [8] Campbell C, Bennett K (2000) A linear programming approach to novelty detection. In: Leen T, Diettrich T, Tresp V (eds) Proceedings of Advances in Neural Information Processing Systems 13 (NIPS'00), MIT Press, Cambridge, MA
- [9] Cook D (2006) Health monitoring and assistance to support aging in place. *Journal of Universal Computer Science* 12(1):15–29
- [10] Egenhofer M, Franzosa R (1991) Point-set topological spatial relations. *International Journal of Geographical Information Systems* 5(2):161–174
- [11] Elman J (1990) Finding structure in time. *Cognitive Science* 14:179–211

- [12] Fine S, Singer Y, Tishby N (1998) The hierarchical hidden markov model: Analysis and applications. *Machine Learning* 32:41–62
- [13] Gopalratnam K, Cook D (2004) Active LeZi: An incremental parsing algorithm for sequential prediction. *International Journal of Artificial Intelligence Tools* 14(1–2):917–930
- [14] Guesgen H (1989) Spatial reasoning based on Allen’s temporal logic. Technical Report TR-89-049, ICSI, Berkeley, California
- [15] Han J, Pei J, Mao R (2004) Mining frequent patterns without candidate generation: A frequent-pattern tree approach. *Data Mining and Knowledge Discovery* 8:53–87
- [16] Heierman E, Cook D (2003) Improving home automation by discovering regularly occurring device usage patterns. In: *Proceedings of the International Conference on Data Mining*
- [17] Jakkula V, Cook D (2008) Anomaly detection using temporal data mining in a smart home environment. *Methods of Information in Medicine*
- [18] Jordan M (ed) (1999) *Learning in Graphical Models*. The MIT Press, Cambridge, MA
- [19] Kalman R (1960) A new approach to linear filtering and prediction problems. *Journal of Basic Engineering* 82:34–45
- [20] Kiecolt-Glaser J, Glaser R (1999) Chronic stress and mortality among older adults. *Journal of the American Medical Association* 282:2259–2260
- [21] Kohonen T (1993) *Self-Organization and Associative Memory*, 3rd edn. Springer, Berlin
- [22] Kohonen T, Oja E (1976) Fast adaptive formation of orthogonalizing filters and associative memory in recurrent networks of neuron-like elements. *Biological Cybernetics* 25:85–95
- [23] Kremer SC (2001) Spatiotemporal connectionist networks: A taxonomy and review. *Neural Computation* 13:249–306
- [24] Magnusson M (2000) Discovering hidden time patterns in behavior: T-patterns and their detection. *Behavior Research Methods, Instruments, and Computers* 32(1):93–110
- [25] Marsland S (2003) Novelty detection in learning systems. *Neural Computing Surveys* 3:157–195
- [26] Marsland S (2009) *Machine Learning: An Algorithmic Perspective*. CRC Press, Boca Raton, Florida
- [27] Marsland S (2009) Using habituation in machine learning. *Neurobiology of Learning and Memory* In press
- [28] Marsland S, Shapiro J, Nehmzow U (2002) A self-organising network that grows when required. *Neural Networks* 15(8-9):1041–1058
- [29] Martin G (1998) *Human Neuropsychology*. Prentice Hall, London, England
- [30] McCarthy J, Hayes P (1969) Some philosophical problems from the standpoint of artificial intelligence. In: Meltzer B, Michie D (eds) *Machine Intelligence*, vol 4, Edinburgh University Press, Edinburgh, Scotland, pp 463–502
- [31] Mozer M (2005) Lessons from an adaptive house. In: Cook D, Das R (eds) *Smart environments: Technologies, protocols, and applications*, pp 273–294

- [32] Mukerjee A, Joe G (1990) A qualitative model for space. In: Proc. AAAI-90, Boston, Massachusetts, pp 721–727
- [33] Newell A, Simon H (1976) Computer science as empirical inquiry: Symbols and search. *Communications of the ACM* 19(3):113–126
- [34] Nguyen NT, Phung DQ, Venkatesh S, Bui H (2005) Learning and detecting activities from movement trajectories using the hierarchical hidden markov model. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp 955–960
- [35] O’Keefe J, Nadel L (1977) *The Hippocampus as a Cognitive Map*. Oxford University Press, Oxford, UK
- [36] Preparata F, Shamos M (1985) *Computational Geometry: An Introduction*. Springer, Berlin, Germany
- [37] Quinn M, Johnson M, Andress E, McGinnis P, Ramesh M (1999) Health characteristics of elderly personal care home residents. *Journal of Advanced Nursing* 30:410–417
- [38] Rabiner L (1989) A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE* 77(2):257–286
- [39] Randell D, Cui Z, Cohn A (1992) A spatial logic based on regions and connection. In: Proc. KR-92, Cambridge, Massachusetts, pp 165–176
- [40] Rankin C, Abrams T, Barry R, Bhatnagar S, Clayton D, Colombo J, Coppola G, Geyer M, Glanzman D, Marsland S, McSweeney F, Wilson D, Wu CF, Thompson R (2009) Habituation revisited: An updated and revised description of the behavioral characteristics of habituation. *Neurobiology of Learning and Memory* In press
- [41] Rissanen J (1989) *Stochastic Complexity in Statistical inquiry*. World Scientific Publishing Company
- [42] Rivera-Illingworth F, Callaghan V, Hagrais H (2007) Detection of normal and novel behaviours in ubiquitous domestic environments. *The Computer Journal*
- [43] Rumelhart D, Hinton G, Williams R (1986) Learning internal representations by error propagation. In: Rumelhart D, McClelland JL (eds) *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, The MIT Press, Cambridge, MA, vol 1, pp 318–362
- [44] Stanley J (1976) Computer simulation of a model of habituation. *Nature* 261:146–148
- [45] Stein RB (1967) The frequency of nerve action potentials generated by applied currents. *Proceedings of the Royal Society, B* 167:64–86
- [46] Sutton CA, Rohanimanesh K, McCallum A (2007) Dynamic conditional random fields: Factorized probabilistic models for labeling and segmenting sequence data. *Journal of Machine Learning Research* 8:693–723
- [47] Sutton R, Barto A (1998) *Reinforcement Learning: An Introduction*. The MIT Press, Cambridge, MA
- [48] Tavenard R, Salah A, Pauwels E (2007) Searching for temporal patterns in ami sensor data. In: Proceedings of AmI2007, pp 53–62
- [49] Thompson R, Spencer W (1966) Habituation: A model phenomenon for the study of neuronal substrates of behavior. *Psychological Review* 73(1):16–43

- [50] Vilain M, Kautz H (1986) Constraint propagation algorithms for temporal reasoning. In: Proc. AAAI-86, Philadelphia, Pennsylvania, pp 377–382