The Thirty-First International Florida
Artificial Intelligence Research Society Conference (FLAIRS-31)

Using Spatio-Temporal Anomalies to Detect Abnormal Behaviour in Smart Homes

Hans W. Guesgen, Dick Whiddett, Inga Hunter, Phoebe Elers, Caroline Lockhart, Amardeep Singh
Massey University,
Palmerston North, New Zealand

Stephen Marsland
Victoria University,
Wellington, New Zealand

Abstract

This paper investigates how spatial and temporal context information can be used in smart homes to detect abnormal behaviours. We discuss how various formalisms, such as probability theory, the Dempster-Shafer calculus, and fuzzy logic, can be used to capture context information and argue that fuzzy logic is the most suitable. We evaluate our approach by analysing one of the CASAS smart home datasets.

Introduction

Most developed countries around the world are facing the problem of an ageing population. Life expectancy is higher than ever before, and people expect to maintain a high-quality, independent lifestyle into retirement and beyond, regardless of infirmity or illnesses such as Alzheimer's or Parkinson's diseases. Unfortunately, this hope is not always met, which means that many elderly people require hospital/residential care at some stage in their life, which places an unsustainable demand on healthcare services (New Zealand Ministry of Health 2016).

Rather than moving an elderly person to a nursing home or employing the continuous support of a caretaker, there is a trend to enable ageing-in-place by utilising smart home technology and the elderly person's social support network (Peek et al. 2014). The smart home uses sensors to monitor the inhabitant's activities in an unobtrusive way, with the aim of detecting behaviours that are potentially harmful and seeking appropriate assistance when necessary. However, without an intelligent reasoning engine, the information obtained from the sensors is of little use. The reasoning engine has to determine which activity is currently taking place and whether this activity is a normal behaviour, or poses a threat to the inhabitant.

Recognising human activities is a challenging task, since they can be quite complex, irregular, and vary substantially between instances. Over the last ten years, a large body of research on activity recognition has been developed. The approaches taken range from logic-based approaches to probabilistic machine learning approaches, e.g., (Augusto and Nugent 2004; Chua, Marsland, and Guesgen 2011; Duong et al. 2005; Gopalratnam and Cook 2004; Lotfi et al. 2012; Mahmoud, Lotfi, and Langensiepen 2014; Rivera-Illingworth, Callaghan, and Hagras 2010; Sánchez, Tentori, and Favela 2007; Tapia, Intille, and Larson 2004; Tong, Chen, and Gao 2015). Although the reported successes are promising, simply determining which activity is taking place is not sufficient to decide whether this particular instance of the activity is abnormal. Abnormality can be evident in things like missing important steps of a behaviour (such as turning the gas off after cooking), but it is often the context that determines problems, in particular the spatial and temporal context (Aztiria et al. 2008; Guesgen and Marsland 2010; Gottfried et al. 2015; Jakkula and Cook 2008; Tavenard, Salah, and Pauwels 2007). For example, it is normal to have meals during the day, but not usual to have a meal at 3am. Or it is normal to sleep in the bedroom rather than the laundry.

In this paper, we discuss how spatio-temporal context information can be represented and how spatio-temporal anomalies can be detected. We argue that using a probabilistic approach has shortfalls, due to incomplete context information, and that using an approach based on the concept of beliefs fails due to its complexity. We then make a case for fuzzy logic, which not only provides a simple and robust mechanism for reasoning about contextual information, but also provides a means to represent imprecise information.

A Case for Fuzzy Logic

Activities often take place in particular contexts (locations, times, and before and after other activities), but there is usually no one-to-one relationship between an activity and the context it occurs in. Rather, given a certain activity, contextual information is determined according to some probability distribution. If A is an activity (e.g., making breakfast) and C some context information (e.g., in the kitchen), then $P(C|A)$ is the probability that C is true if A occurs (e.g., the activity takes place in the kitchen if we know that the activity is making breakfast).

Given the conditional probability $P(C|A)$, we can calculate the conditional probability $P(A|C)$ using Bayes' rule:

$$P(A|C) = \frac{P(C|A) \cdot P(A)}{P(C)}$$

. There are several difficulties with using probability to predict activies. A large amount of data is needed, since we need to know the conditional probability for each context and activity. Moreover, we not only need $P(A|C)$ but also $P(A|C_1, \ldots, C_n)$, (where there are n different pieces of contextual information), since it is often correlated. Second, context data are not necessarily discrete. It is common to wish to turn quantise data, for example turning precise times into times of day (for example, 1515 could be afternoon). However, putting hard thresholds on such binning is problematic.

(Clermont 2012) raises an interesting point when considering probabilistic reasoning in law, which is that combining probabilities leads to counter-intuitive results, the so-called conjunction paradox. Although not directly related to the problem addressed in this paper, his example illustrates the challenge we are facing:

> We purport to decide civil cases according to a more-probable-than-not standard of proof. We would expect this standard to take into account the rule of conjunction, which states that the probability of two independent events occurring together is the product of the probability of each event occurring separately. The rule of conjunction dictates that in a case comprised of two independent elements the plaintiff must prove each element to a much greater degree than 50%: only then will the plaintiff have shown that the probability that the two elements occurred together exceeds 50%. Suppose, for example, that a plaintiff must prove both causation and fault and that these two elements are independent. If the plaintiff shows that causation is 60% probable and fault is 60% probable, then he apparently would have failed to satisfy the civil standard of proof because the probability that the defendant both acted negligently and caused injury is only 36%.

Replacing ‘causation’ and ‘fault’ with two different contextual observations shows the difficulty. In fact, this difficulty arises from the fact that total probability is always 100%. Another problem that this raises is that when our knowledge is incomplete, so that we do not have prior information about something, when a probability of the unknown event has to be estimated, this is often based on nothing more than guesswork, as otherwise it will be assigned probability 0, meaning that it is impossible.

Dempster-Shafer theory (Dempster 1967; Shafer 1976) offers a way out of this dilemma. It uses the concepts of belief (Bel) and plausibility (Pl) instead of probability to formulate uncertainty, where classical probability lies between belief and plausibility:

$$Bel(A|C) \leq P(A|C) \leq Pl(A|C)$$

Belief and plausibility are defined in Dempster-Shafer theory on the basis of a mass function, which assigns basic probabilities to the power set of the so-called frame of discernment:

$$m : 2^U \rightarrow [0, 1]$$

In the context of activity recognition, the frame of discernment is the set of possible activities that can be observed by the ambient intelligence of the smart home. The mass assignments to the activities must add up to 1, and the mass assignment of the empty set has to be 0:

$$\sum_{X \in U} m(X) = 1 \qquad m(\emptyset) = 0$$

For example, if we were able to observe only three types of activity, having breakfast (A_b), eating dinner (A_d), and washing dishes (A_c), the frame of discernment would be $U = \{A_b, A_c, A_d\}$. Further, if we knew that the activity that just took place occurred in the kitchen and that 80% of kitchen activities are either having breakfast or washing dishes, then we would assign 0.8 as basic probability to the set $\{A_b, A_c\}$:

$$m_s(\{A_b, A_c\}) = 0.8$$

If we had no other spatial context information, we would assign the remaining probability to the whole frame of discernment:

$$m_s(\{A_b, A_c, A_d\}) = 0.2$$

If we have two mass functions, for instance one for some spatial context information and the other for some temporal context information, we can combine these for every nonempty subset X of the frame of discernment in the following way:

$$[m_s \oplus m_t](X) = \frac{\sum_{V \cap W = X} m_s(V) \cdot m_t(W)}{1 - \sum_{V \cap W = \emptyset} m_s(V) \cdot m_t(W)}$$

The combined mass assignment for the empty set is 0:

$$[m_s \oplus m_t](\emptyset) = 0$$

This assumes that the weight of conflict is not equal to 1:

$$\sum_{V \cap W = \emptyset} m_s(V) \cdot m_t(W) \neq 1$$

To illustrate the combination of mass functions, assume that m_s makes mass assignments as in the example above. Further assume that m_t assigns 0.6 to the set $\{A_c, A_d\}$, which means that the observed activity is either washing dishes or eating dinner with a certainty of 60% (since the activity took place in the evening and 60% of evening events are either washing dishes or eating dinner). If we had no other temporal context information, we would assign the remaining probability to the whole frame of discernment:

$$m_t(\{A_c, A_d\}) = 0.6 \qquad m_t(\{A_b, A_c, A_d\}) = 0.4$$

Combining the spatial and temporal evidence results in the mass assignment shown in Table 1, where m_c denotes the combined mass assignment function.

Although this approach solves the problem of evidence not adding up to 100% (which is not the same as mass assignments adding up to 100%, since there is an explicit ‘don’t know’ category), it is not necessarily a practical solution to reasoning about context information in general, since Dempster-Shafer theory is an even more complex framework than probability theory, even if it is used in a qualitative form as discussed in (Guesgen and Marsland 2015). There is also the question of how to obtain the mass assignments. One way is to base them on smart home datasets (see next section). Given a particular context, whenever that

$\oplus$	$m_s(\{A_b, A_c\}) = 0.8$	$m_s(\{A_b, A_c, A_d\}) = 0.2$
$m_t(\{A_c, A_d\}) = 0.6$	$m_c(\{A_c\}) = 0.48$	$m_c(\{A_c, A_d\}) = 0.12$
$m_t(\{A_b, A_c, A_d\}) = 0.4$	$m_c(\{A_b, A_c\}) = 0.32$	$m_c(\{A_b, A_c, A_d\}) = 0.08$

Table 1: Example of combining spatial and temporal mass functions in Dempster-Shafer theory.

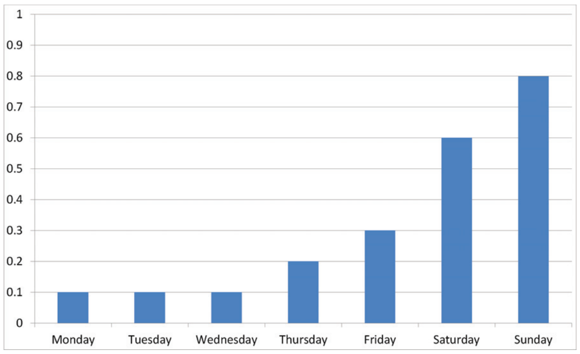

Figure 1: Graphical representation of a membership functions that determines the degree to which a restaurant visit falls on a particular day of the week.

context occurs, we increment a counter for the set of activities happening in this context. For example, each evening we memorise the activities occurring on that evening. If the only behaviour is eating dinner, then the counter for $\{A_d\}$ is incremented by 1, if the only activity is washing dishes, then the one for $\{A_c\}$ is incremented, and if both of them occur, the counter for $\{A_c, A_d\}$ is incremented. After that, the normalised counters are used as mass assignments.

Context Fuzzy Sets

Unlike traditional sets, fuzzy sets allow their elements to belong to the set with a certain degree. Rather than deciding whether an element d does or does not belong to a subset A of a domain D, we determine for each element of D the degree with which it belongs to the fuzzy set $\tilde{A}$. In other words, a fuzzy subset $\tilde{A}$ of a domain D is a set of ordered pairs, $(d, \mu_{\tilde{A}}(d))$, where $d \in D$ and $\mu_{\tilde{A}} : D \to [0, 1]$ is the membership function of $\tilde{A}$. The membership function replaces the characteristic function of a classical subset $A \subset D$.

Rather than asking the question of what is the probability of a certain activity occurring in a particular context, we now pose the question as follows. Given some context information C, to which degree is a particular activity a C-activity. For example, if C is the day of the week, then we can ask for the degree of the activity to be a Monday activity, Tuesday activity, and so on. In terms of fuzzy sets, we define D as the set of the seven days of the week and $\mu_{\tilde{A}}$ as the membership function that determines to what degree the activity occurs on a particular day. The membership function of such a fuzzy set is shown graphically in Figure 1. Unlike probabilities, the membership grades do not need to add up to one.

In the example above, the context information is still crisp

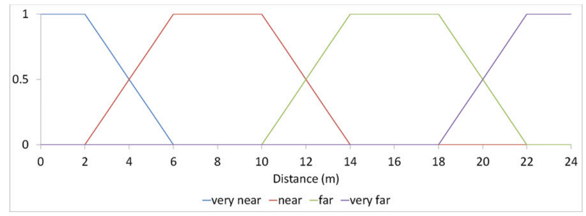

Figure 2: A fuzzy set that maps distances to the qualitative values very near, near, far, and very far.

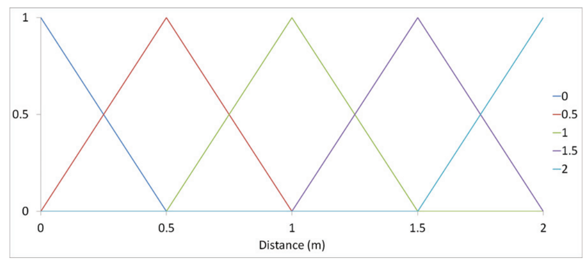

Figure 3: A fuzzy set that approximates distances with a granularity of half a metre.

information, despite the fact that it is used in a fuzzy set: for any restaurant visit, we can determine precisely which day of the week it occurred on. Other context information might not be precise; for example, we may know that an activity occurs near the kitchen but the precise distance (e.g., in metres) may well be unknown. In this case, we can represent the context information itself as a fuzzy set, as illustrated in Figure 2.

Similarly, we can define a fuzzy set that expresses distances by rounding them to the closest half metre – something we as humans often do when we perceive distances, although not necessarily always on the same scale (see Figure 3).

All of the examples above relate particular activities to either spatial or temporal context. There is another form of spatio-temporal context awareness, which looks at two activities and captures how they are related to each other in space and/or time. For example, if the current activity is having breakfast, then activities such as leaving the house (to go to work) or having another cup of coffee (at the weekend) may be quite likely, whereas going to bed is not. This type of information can also be captured in fuzzy sets.

We have evaluated our approach using the CASAS datasets (Cook 2012) [1]. These datasets cover a wide range of

[1] Available at http://www.ailab.wsu.edu/casas/datasets/

scenarios, from single-resident apartments to multi-resident homes with interweaved activities of daily living. The datasets are primarily based on motion sensors, but also include environmental sensors, object sensors, and wearable sensors. Some of the datasets are fully annotated (i.e., contain ground truth), while others are only partially annotated or not annotated at all. The particular dataset we have used is the Aruba dataset, which is the annotated data of a single person living their daily life.

Table 2 shows our analysis of the Aruba dataset. Each entry in the table shows the (normalised) frequency with which a particular activity is followed by another activity. We interpret these frequencies as fuzzy membership grades, which means that we can create a successor fuzzy set for each activity present in the dataset.

For each activity, the entries in the table differ significantly for the successor activities, which means the resulting fuzzy sets can inform the activity recognition process effectively in a number of ways. For example, we can select for each activity the three successors with the largest membership grades and bias the activity recogniser towards these successors. Figure 4 visualises this boosting process using a successor graph.

As the examples have shown, fuzzy sets can be used for associating activities with context information and for representing imprecise context information. Fuzzy set theory also provides us with a means to convert fuzzy sets back to crisp sets, which is achieved with the notion of an α-level set. Let $\tilde{A}$ be a fuzzy subset in D, then the (crisp) set of elements that belong to the fuzzy set $\tilde{A}$ with a membership grade of at least α is called the α-level set of $\tilde{A}$:

$$A_\alpha = d \in D | \mu_{\tilde{A}}(d) \geq \alpha$$

If the membership grade is strictly greater than α, then the set is referred to as a strong α-level set.

Reasoning with Context Fuzzy Sets

To avoid the problem of accumulating values, which we encounter in probability theory and Dempster-Shafer theory, we choose one of the schemes for combining fuzzy sets that was proposed in (Zadeh 1965). Given two fuzzy sets $\tilde{A}_1$ and $\tilde{A}_2$ with membership functions $\mu_{\tilde{A}_1}$ and $\mu_{\tilde{A}_2}$, respectively, then the membership function of the intersection $\tilde{A}_3 = \tilde{A}_1 \cap \tilde{A}_2$ is pointwise defined by:

$$\mu_{\tilde{A}_3}(d) = \min\{\mu_{\tilde{A}_1}(d), \mu_{\tilde{A}_2}(d)\}$$

Analogously, the membership function of the union $\tilde{A}_3 = \tilde{A}_1 \cup \tilde{A}_2$ is pointwise defined by:

$$\mu_{\tilde{A}_3}(d) = \max\{\mu_{\tilde{A}_1}(d), \mu_{\tilde{A}_2}(d)\}$$

The membership grade for the complement of a fuzzy set $\tilde{A}$, denoted as $\neg\tilde{A}$, is defined in the same way as the complement in probability theory:

$$\mu_{\neg\tilde{A}}(d) = 1 - \mu_{\tilde{A}}(d)$$

(Zadeh 1965) stresses that the min/max combination scheme is not the only scheme for defining intersection and union of fuzzy sets, and that it depends on the context which scheme is the most appropriate. While some of the schemes are based on empirical investigations, others are the result of theoretical considerations (Dubois and Prade 1980; Klir and Folger 1988). However, Nguyen et al. (Nguyen, Kreinovich, and Tolbert 1993) proved that the min/max operations are the most robust operations for combining fuzzy sets, where robustness is defined in terms of how much impact uncertainty in the input has on the error in the output.

Figure 5 shows the overall architecture of the activity recognition system. Sensors positioned throughout the smart home, such as motion sensors, temperature sensors, door sensors, etc. produce a data stream that feeds into the activity recogniser. The activity recogniser first produces candidates of activities that might correspond to the observed data stream. By using spatio-temporal fuzzy sets, it then decides whether the activity is normal or not.

Conclusion

This paper discusses how context information can be represented and how it can be used for activity recognition in smart homes. Since context information is often incomplete, imprecise, and complex, probabilistic approaches are impractical and therefore potentially unhelpful in real-world applications. Fuzzy logic offers a way around this, as it provides robust mechanisms for dealing with uncertainty.

Although we used context information in this paper to boost the activity recognition process, this is not its only application. Just determining which activity takes place in the smart home is usually not sufficient, because we often also need to know whether an activity is normal or abnormal. One way to distinguish between a normal and an abnormal activity is to look at the context in which it takes place. If the context is different from the usually observed context, the activity might be abnormal.

Acknowledgement

A preliminary version of this paper was presented at the AAAI-13 Workshop on Space, Time, and Ambient Intelligence in Bellevue, Washington. The authors are grateful for the valuable comments and suggestions made by the participants of the workshop. This work has been supported by Explorer Grant 16/679 from the Health Research Council of New Zealand.

References

Augusto, J., and Nugent, C. 2004. The use of temporal reasoning and management of complex events in smart homes. In *Proc. ECAI-04*, 778–782.

Aztiria, A.; Augusto, J.; Izaguirre, A.; and Cook, D. 2008. Learning accurate temporal relations from user actions in intelligent environments. In *Proc. 3rd Symposium of Ubiquitous Computing and Ambient Intelligence*, 274–283.

Chua, S.-L.; Marsland, S.; and Guesgen, H. 2011. Unsupervised learning of patterns in data streams using compression and edit distance. In *Proc. IJCAI*, 1231–1236.

	Successor											
Activity	Bed to toilet	Eating	Enter home	House-keeping	Leave home	Meal preparation	Presumed sleeping	Relax	Resperate	Sleeping	Wash dishes	Work
Bed to toilet		0.000	0.000	0.000	0.000	0.006	0.025	0.000	0.000	0.967	0.000	0.000
Eating	0.000		0.000	0.004	0.207	0.176	0.000	0.380	0.000	0.008	0.172	0.048
Enter home	0.000	0.011		0.004	0.238	0.375	0.002	0.317	0.000	0.002	0.000	0.046
Housekeeping	0.000	0.000	0.000		0.363	0.121	0.000	0.393	0.030	0.000	0.000	0.090
Leave home	0.000	0.000	1.000	0.000		0.000	0.000	0.000	0.000	0.000	0.000	0.000
Meal preparation	0.000	0.198	0.000	0.003	0.053		0.000	0.718	0.000	0.000	0.000	0.024
Presumed sleeping	0.222	0.000	0.000	0.000	0.055	0.722		0.000	0.000	0.000	0.000	0.000
Relax	0.000	0.016	0.000	0.017	0.166	0.505	0.011		0.002	0.194	0.021	0.063
Resperate	0.000	0.000	0.000	0.000	0.166	0.333	0.000	0.500		0.000	0.000	0.000
Sleeping	0.411	0.000	0.000	0.000	0.000	0.539	0.002	0.046	0.000		0.000	0.000
Wash dishes	0.000	0.015	0.000	0.046	0.093	0.062	0.000	0.703	0.000	0.000		0.078
Work	0.000	0.015	0.000	0.030	0.203	0.300	0.000	0.398	0.007	0.030	0.015	

Table 2: Normalised frequencies with which activities follow other activities. An entry in row i and column j shows the frequency with which activity j follows activity i.

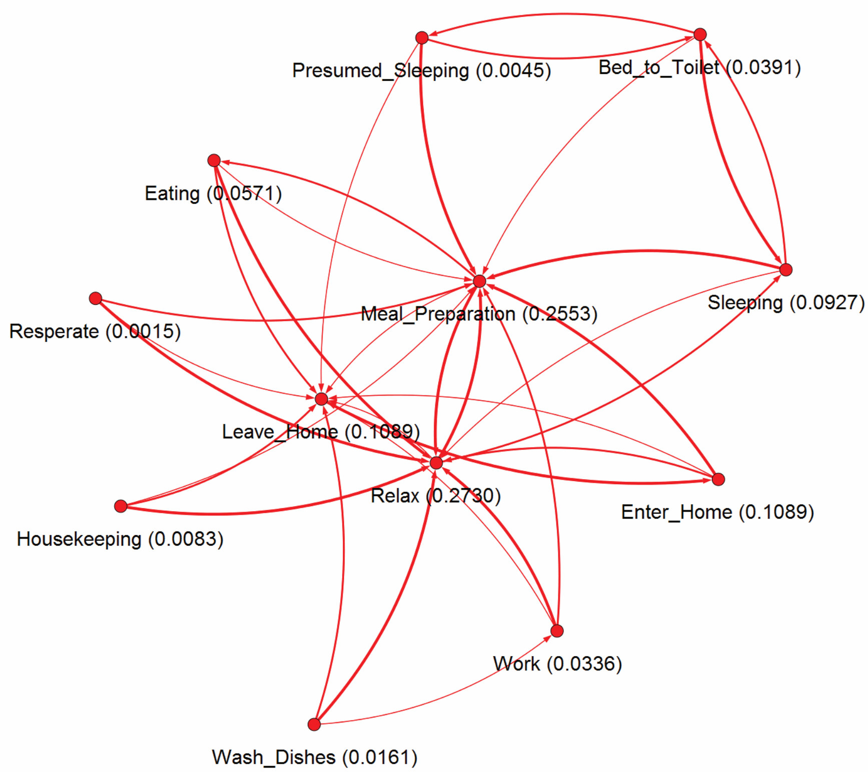

Figure 4: Graph showing the three successors with the highest fuzzy membership grades for each activity. The successor activity with the highest grade is connected to its predecessor by a thick arrow, the one with the second highest by a medium arrow, and the activity with the third highest by a thin arrow.

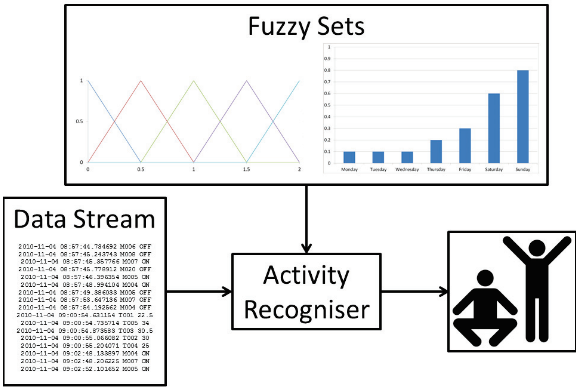

Figure 5: System architecture of an activity recogniser. The recogniser is triggered by a data stream resulting from the sensors installed in the smart home. Spatial and temporal information in the form of fuzzy sets is used to decide whether a recognised activity is normal or not.

Clermont, K. 2012. Death of paradox: The killer logic beneath the standards of proof. Cornell Law Faculty Working Papers 94, Cornell University, Ithaca, New York.

Cook, D. 2012. Learning setting-generalized activity models for smart spaces. *Intelligent Systems* 27(1):32–38.

Dempster, A. 1967. Upper and lower probabilities induced by a multivalued mapping. *The Annals of Mathematical Statistics* 38(2):325–339.

Dubois, D., and Prade, H. 1980. *Fuzzy Sets and Systems: Theory and Applications*. London, England: Academic Press.

Duong, T.; Bui, H.; Phung, D.; and Venkatesh, S. 2005. Activity recognition and abnormality detection with the switching hidden semi-Markov model. In *Proc. CVPR-05*, 838–845.

Gopalratnam, K., and Cook, D. 2004. Active LeZi: An incremental parsing algorithm for sequential prediction. *International Journal of Artificial Intelligence Tools* 14(1–2):917–930.

Gottfried, B.; Aghajan, H.; Wong, K.-Y.; Augusto, J.; Guesgen, H.; Kirste, T.; and Lawo, M. 2015. Spatial health systems: When humans move around. In Holzinger, A.; Röcker, C.; and Ziefle, M., eds., *Smart Health: Open Problems and Future Challenges*, Lecture Notes in Computer Science 8700. Cham, Switzerland: Springer International Publishing. 41–69.

Guesgen, H., and Marsland, S. 2010. Spatio-temporal reasoning and context awareness. In Nakashima, H.; Aghajan, H.; and Augusto, J., eds., *Handbook of Ambient Intelligence and Smart Environments*. Berlin, Germany: Springer. 609–634.

Guesgen, H., and Marsland, S. 2015. Human behaviour recognition in ambient intelligent environments. In Curran, K., ed., *Recent Advances in Ambient Intelligence and Context-Aware Computing*. Hershey, Pennsylvania: IGI Global. 107–121.

Jakkula, V., and Cook, D. 2008. Anomaly detection using temporal data mining in a smart home environment. *Methods of Information in Medicine* 47(1):70–75.

Klir, G., and Folger, T. 1988. *Fuzzy Sets, Uncertainty, and Information*. Englewood Cliffs, New Jersey: Prentice Hall.

Lotfi, A.; Langensiepen, C.; Mahmoud, S.; and Akhlaghinia, M. 2012. Smart homes for the elderly dementia sufferers: Identification and prediction of abnormal behaviour. *Ambient Intelligence and Humanized Computing* 3:205–218.

Mahmoud, S.; Lotfi, A.; and Langensiepen, C. 2014. User activities outliers detection; integration of statistical and computational intelligence techniques. *Computational Intelligence*.

New Zealand Ministry of Health. 2016. Health of older people strategy: Consultation draft. https://www.health.govt.nz/system/files/documents/publications/health-older-people-strategy-consultation-draft-july2016-v2.pdf. Accessed: 2018-01-12.

Nguyen, H.; Kreinovich, V.; and Tolbert, D. 1993. On robustness of fuzzy logics. In *Proc. 2nd IEEE International Conference on Fuzzy Systems*, 543–547.

Peek, S.; Wouters, E.; van Hoof, J.; Luijkx, K. G.; Boeije, H.; and Vrijhoef, H. 2014. Factors influencing acceptance of technology for aging in place: A systematic review. *International Journal of Medical Informatics* 83(4):235–248.

Rivera-Illingworth, F.; Callaghan, V.; and Hagras, H. 2010. Detection of normal and novel behaviours in ubiquitous domestic environments. *The Computer Journal* 53(2):142–151.

Sánchez, D.; Tentori, M.; and Favela, J. 2007. Hidden markov models for activity recognition in ambient intelligence environment. In *Proc. Mexican International Conference on Current Trends in Computer Science*, 33–40.

Shafer, G. 1976. *A Mathematical Theory of Evidence*. Princeton, New Jersey: Princeton University Press.

Tapia, E.; Intille, S.; and Larson, K. 2004. Activity recognition in the home using simple and ubiquitous sensors. In *Proc. PERVASIVE-04*, 158–175.

Tavenard, R.; Salah, A.; and Pauwels, E. 2007. Searching for temporal patterns in ami sensor data. In *Proc. of AmI2007*, 53–62.

Tong, Y.; Chen, R.; and Gao, J. 2015. Hidden state conditional random field for abnormal activity recognition in smart homes. *Entropy* 17:1358–1378.

Zadeh, L. 1965. Fuzzy sets. *Information and Control* 8:338–353.