Spatio-Temporal Footprints

Hans W. Guesgen, Massey University, New Zealand
Stephen Marsland, Massey University, New Zealand

Abstract

The recognition of human behaviour from sensor observations is an important area of research in smart homes and ambient intelligence. In this paper, we introduce the idea of spatio-temporal footprints, which are local patterns in space and time that should be similar across repeated occurrences of the same behaviour. We discuss the spatial and temporal mapping requirements of these footprints, together with how they may be used.

Keywords: Abnormality Detection; Behaviour Recognition; Spatio-Temporal Reasoning; Space-Time Invariants

1 Introduction

A common task that an ambient intelligence system could be required to perform is recognising human behaviour from observations in the environment; this can be useful for a variety of applications from monitoring the activities of elderly patients to identifying appropriate lighting and heating conditions (Cook, 2006; Mozer, 2005). The observations on which such recognition is based can range from direct observations made by video cameras to indirect observations detected by sensors. Although video cameras give a more complete picture, and hence might lend themselves more easily to recognising behaviours (with a consequent increase in the amount of computational processing required), it is often behaviour recognition based on sensors that is the preferred option, since the latter is less obtrusive and therefore more easily accepted in applications such as smart homes.

There is a significant body of research on behaviour recognition based on sensor data, which ranges from logic-based approaches to probabilistic machine learning approaches (Augusto & Nugent, 2004; Chua et al., 2009; Duong et al., 2005; Gopalratnam & Cook, 2004; Rivera-Illingworth et al., 2007; Tapia et al., 2004). Although the reported successes are promising, it has become clear that all approaches fall short of being perfect. Due to the limited information that is in the sensor data, noise, and the inherently complexity of human behaviours, it is often impossible to determine the correct behaviour from the sensor data alone, in particular if behaviours are overlapping or are being executed by more than one person.

Several researchers have realised that additional information can be useful to boost the behaviour recognition process (Aztiria et al., 2008; Jakkula & Cook, 2008; Tavenard et al., 2007). In this article, we focus on how spatio-temporal information, enriched with context information, can be used for this purpose. When a particular activity occurs, like preparing breakfast, it leaves a ‘footprint’ in space-time, i.e., a particular pattern of sensor observations in some set of locations over some period of time. The activity starts at some specific time and in some specific location, goes on for a specific duration in some specific area, and terminates at a specific time at some specific location. Since footprints differ from behaviour to behaviour—but often relatively little between different instances of the same behaviour—we can use these to inform the behaviour recognition process: if something is happening at 07:00 in the kitchen, it is more likely to be preparing breakfast than taking a shower. We can also use them to detect abnormal behaviour: if the inhabitant of the smart home uses the shower at 03:00 (when usually this is not a footprint that is seen), then this can be interpreted as abnormal.

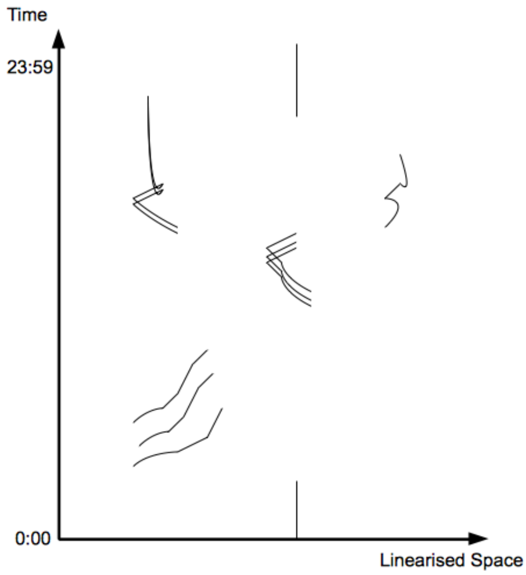

Figure 1. An example of space-time with footprints.

Figure 1 shows an example of a possible set of footprints over three days, with a linear time axis that repeats each day, and a single space axis that could identify rooms, or similar (this is discussed in more detail in the next section). It can be seen that some behaviours repeat more-or-less identically over the three days, while others only occur once. The challenge with such representation of behaviour as footprints is to identify and recognise the various behaviours that are represented.

The rest of this article investigates the role of spatio-temporal footprints in more detail. We start with a discussion of what space-time means in the context of smart homes, arguing that there is more than one space-time (or more precisely, representation of space-time). We then look at how behaviours leave footprints in space-time and explore invariants in these footprints, with the goal of classifying different forms of invariants. Finally we will look into how the footprints are distributed in space-time and what influences this distribution.

2 Space-Time

When reasoning about time, we usually associate a time axis with the data. The time axis might use a calendar as reference system and absolute dates/times to refer to points on the axis. Or it might use some artificial start point as zero time, such as the time when the smart home became operational, and some counter to advance time along the time axis. In the latter case, we would not be able to refer back to times before the birth of the smart home, while the first case would provide an infinite extension of time into both the past and the future.

As we will see later when discussing footprint invariants, it may sometimes be advantageous to view the time axis as a circular reoccurrence of time points. For example, if we are only interested in when behaviours occur during the day, then we might want to abstract from years, months, and days, which would leave us with references to times of the day. At the end of the day, we would ‘warp’ time and start at the beginning of the time axis again.

Independently of whether we use linear or circular time, we still have to decide whether we view time as continuous or discrete. For example, when referring to 13:00, do we really mean this exact point on a continuous time axis? If this is the case, then a behaviour that occurs one millisecond after this time point would not match 13:00, unless we allow for ‘fuzzy’ matches. On the other hand, if we view it as a discrete time stamp, surrounded by, say, 12:55 and 13:05, then it would make sense to associate the behaviour with 13:00 rather than 12:55 or 13:05. This can be thought of as temporal ‘resolution’.

Similar considerations can be made when referring to space. Although space is more complex than time (partially because we can move freely in space, but not in time), it has many similarities to time. In the simplest case, it has the same dimensionality, for example if we can move only along a predefined trajectory (Mukerjee & Joe, 1990). The trajectory can be viewed as continuous, in which case we would associate the distance from the origin with locations on the spatial axis, or it can be discrete, in which case we need a way to associate distances with reference points on the spatial axis.

As with time, we can envision different representations of space. Not only can we extend the dimensionality of space to two or three, we can also move away from a canonical Euclidean space to a more abstract space. For example, we can use the rooms of the smart home to define space points, or the areas covered by the sensors of the smart home. Given some knowledge of the physical locations of sensors, the house can infer that when sensor events occur, the house inhabitant must be in the physical vicinity of that sensor (obviously, there are exceptions to this for certain sensors such as thermometers, and for remote-controlled or time-controlled devices). The spatial pattern could then be some mapping between sensor locations, which could be based on the underlying physical layout of the house, but does not have to be.

The importance of this is that different resolutions can make recognition of particular footprints easier or more difficult. As the resolution becomes finer, footprints that appear to be exactly aligned start to separate, meaning that identifying them as examples of the same behaviour can require a clustering algorithm. Breakfast might not occur exactly at 08:00 each day, but it is likely that it happens at times that we associate with morning. Or reading a book might always occur in the lounge, but sometimes while sitting on the sofa and sometimes while lying back in an easy chair.

3 Footprint Identification

Each activity leaves a footprint in space-time (see Figure 1). If it is possible to determine mappings between footprints for the different instances of the same behaviour occurring in the smart home, then we can use that information to improve the behaviour recognition process. This section discusses what type of mappings we might usefully want to identify. There are effectively three different things that we might want to detect (where the first two are positive—examples of the same pattern—while the final one is negative):

- The same (or very similar) pattern occurring at the same time and place
- The same (or very similar) pattern occurring at different times and/or places
- Different patterns occurring at the same time and place as another

As an example of the first of these, consider the case that breakfast always takes place in the kitchen at 08:00. Then the mapping between this footprint on different days is simply the identity, as it leaves the same footprint in space-time each time it occurs. Of course, this assumes that we have chosen a suitable representation for space-time (e.g., one that only looks at the time of the day in a discrete way and uses the rooms of the smart home as spatial entities). It is unlikely that we observe this invariant very often, as it would require very rigid patterns of behaviour and a relatively abstract form of space-time.

The next type of transformation is that where we match patterns despite shifts in space-time. Here we can distinguish among three different types of shift:

- Shifts on the time axis only
- Shifts on the space axes
- Shifts on all axes

For example, the afternoon tea break might occur at different times but always in the lounge (time-only shift), while the afternoon nap might occur always at the same time but either in the bedroom or in the lounge (space-only). An example of a (coupled) shift in space-time would be breakfast at the weekend: while breakfast during the week happens at 08:00 in the kitchen, it might be shifted to 10:00 in the dining room on Saturday and Sunday.

The particular footprint that identifies a behaviour is caused by some set of sensors being activated in time. While we want some robustness to minor variation in the footprint, we need to be careful that just because two footprints occur in the same place and time, they are not necessarily the same event. To continue the breakfast analogy, a person working at home could go into the kitchen at 10:00 to make a cup of tea. This is not an example of them having a second breakfast (the weekend one), but a different behaviour that happens to occur at the same time and place.

There is one example of footprint change that we might need to be particularly careful about, which is the deformation of footprints in space-time. This again might be restricted to particular axes such as time, or might include all axes. We start the discussion of this by restricting ourselves to the time axis. The projection of a footprint onto the time axis is a time interval, consisting of a start point and an end point. We can restrict the deformation of the time interval to one of its boundaries, either the start or the end point, but not both. An example would be breakfast that always starts at 08:00 but might last between 10 and 20 minutes, or doing the dishes after dinner might start at any time, but always finishes in time to be ready for the soap opera on TV at 18:00. If we do not restrict ourselves to one of the boundaries, we obtain a deformation which (more or less) keeps the centre of the interval invariant. For example, taking a shower might always occur at around 20:00, starting about 10 minutes before that time and ending about 10 minutes after that time.

This same effect can be seen for space. In this case, we have to consider boundaries of regions in (usually) multi-dimensional space rather than start and end points of time intervals, which means that instead of considering two classes of deformation (one boundary vs. two boundaries), we have to consider an infinite number of classes. One way to achieve this is by determining the percentage of the boundary that stays the same and associating the behaviours with a finite set of classes that are given by a range of percentages. For example, we might want to distinguish just between those deformations that change more than 50% of the boundary and those that keep at least 50% of the boundary invariant.

It should be noted that our discussion of footprint invariants is closely related to the discussion around neighbourhood graphs in (Freksa, 1992), where Freksa introduces three forms of neighbourhood graphs for Allen's temporal logic (Allen, 1983). The graphs are based on shifts and two forms of deformations of time interval, in a way closely related to the one outlined above. Similar discussions can be found around the region connection calculus, which is used for reasoning about spatial relations (Randell et al., 1992).

Thus, we are lead to consider two related methods of footprint identification: identifying the same footprint occurring in different places in space-time, and distinguishing between different footprints that occur at the same place in space-time.

Footprint Probability Distributions

The examples of the previous sections indicate that footprints are not necessarily distributed evenly in space-time. A straightforward way to find the distribution of footprints is to empirically approximate for each behaviour the probability of a footprint occurring at a particular location (or within a particular cluster) in space-time. This does not require any extra knowledge about the behaviour or about particular regions in space-time, but it does require enough data to approximate the probabilities within reasonable error margins.

As an alternative to this approach, we can analyse the regions occupied by a particular behaviour in space-time in order to find the distributions of the corresponding footprints. Breakfast on weekends as opposed to weekdays is an example for that, as illustrated previously. Obviously, the footprints are not distributed evenly over the two clusters, but have a higher density in the weekday breakfast cluster than in the weekend breakfast cluster. We know that there are five weekdays per week, but only two weekend days (for the average working person). Assuming that the breakfast behaviour occurs exactly once per day, we conclude that the probability of a breakfast footprint being in the weekend cluster of breakfast footprints is 2/7, whereas the probability of it being in the weekday one is 5/7. Taking this approach a step further, we can then compute conditional probabilities, which give us further insides into where a footprint is located in space-time. For example, if we know that a behaviour occurred on a weekday, then the conditional probability of it being in the weekday cluster is 1 (and 0 for the weekend cluster).

In general, this leads to an approach where context information is taken into consideration when behaviours are related to space-time footprints. In the example above, the context is of a spatio-temporal nature, but this does not have to be the case, as it can make sense to utilise other types of context information as well, such as:

- **Linked behaviours.** If a person has already had breakfast then they are unlikely to be having a second one, and if they have just had a shower they are unlikely to be having a bath. This kind of data can help to separate out the different footprints that might be recognised at the current time.
- **Environmental information.** If it is cold outside and not all rooms of the home are heated properly, then the footprints of certain behaviours might shift in space-time along the spatial axes. Rather than taking a meal in the dining room, the inhabitant might choose to have it in the lounge where there is a fireplace.
- **Personal information.** If the inhabitant is sick, he or she might choose to go to bed earlier than usual. This most likely has an effect on the footprints of events happening towards the end of the day, which would shift along the temporal axis.
- **Socio-economic information.** If there is a recession, the inhabitant might choose to save costs and therefore might decide to reduce the duration of hot showers. As a consequence, the footprint of that behaviour would be deformed.

Although in principle there is no limit to how much context information we use to get a better understanding of the relationship between behaviours and their space-time footprints, it is not practical to use context information excessively. Each bit of information requires us to explicitly model the correlation between the information and its impact on the behaviour–footprint relation, which requires a significant amount of world knowledge. In other words, we trade off the need for sufficient training data against the need for explicit modelling.

Conclusion

We have presented a representation of behaviours as patterns of activity in space-time, where each behaviour is represented as a trajectory of sensor events over some relatively short time window. It is hoped that by representing the behaviours in this way, individual behaviours will be more clearly recognised despite translation in space, time, or both. Additionally, it may give a useful pictorial representation of events that enables a carer to analyse the actions of a person and identify abnormal events once the smart home has raised an alarm.

Some of the challenges of recognising and using such footprints are caused by the natural variability between different instances of the same behaviour. We have discussed this in the context of things taking slightly more or less time, and the location changing, but it is also the case for different orderings of the actions within a behaviour. The extent to which this is a problem will have to be examined once we are able to use real data to examine the footprint pattern.

References

Allen, J. (1983). Maintaining knowledge about temporal intervals. *Communications of the ACM, 26*, 832-843.

Augusto, J., & Nugent, C. (2004). The use of temporal reasoning and management of complex events in smart homes. In *Proceedings of ECAI-04,* Valencia, Spain (pp. 778-782) .

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

Chua, S.-L., Marsland, S., & Guesgen, H. (2009). Spatio-temporal and context reasoning in smart homes. In *Proceedings of the COSIT-09 Workshop on Spatial and Temporal Reasoning for Ambient Intelligence Systems,* Aber Wrac'h, France (pp. 9-20).

Cook, D. (2006). Health monitoring and assistance to support aging in place. *Journal of Universal Computer Science, 12*(1), 15-29.

Duong, T., Bui, H., Phung, D., & Venkatesh, S. (2005). Activity recognition and abnormality detection with the switching hidden semi-Markov model. In *Proceedings of CVPR-05* (pp. 838-845).

Freksa, C. (1992). Temporal reasoning based on semi-intervals. *Artificial Intelligence, 54*, 199-227.

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.

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

Mozer, M. (2005). Lessons from an adaptive house. In D. Cook & R. Das (Eds.), *Smart environments: Technologies, protocols, and applications* (pp. 273-294).

Mukerjee, A., & Joe, G. (1990). A qualitative model for space. In *Proceedings of AAAI-90,* Boston (pp. 721-727).

Randell, D., Cui, Z., & Cohn, A. (1992). A spatial logic based on regions and connection. In *Proceedings of KR-92,* Cambridge, MA (pp. 165-176).

Rivera-Illingworth, F., Callaghan, V., & Hagras, H. (2007). Detection of normal and novel behaviours in ubiquitous domestic environments. *The Computer Journal*.

Tapia, E., Intille, S., & Larson, K. (2004). Activity recognition in the home using simple and ubiquitous sensors. In *Proceedings of PERVASIVE-04,* Vienna, Austria (pp. 158-175).

Tavenard, R., Salah, A., & Pauwels, E. (2007). Searching for temporal patterns in ami sensor data. In *Proceedings of AmI2007* (pp. 53-62).

Author Biographies

Hans Guesgen is a professor of computer science in the School of Engineering and Advanced Technology at Massey University in Palmerston North, New Zealand. He holds a diploma in computer science and mathematics of the University of Bonn, a doctorate in computer science of the University of Kaiserslautern, and a higher doctorate (Habilitation) in computer science of the University of Hamburg, Germany. He worked as a research scientist at the German National Research Center of Computer Science (GMD) at Sankt Augustin from 1983 to 1992. During this period he held a one-year post-doctoral fellowship at the International Computer Science Institute in Berkeley, California. In 1992 he joined the Computer Science Department of the University of Auckland, where he worked until moving to Massey University in 2007. His research interests include ambient intelligence and spatio-temporal reasoning.

Stephen Marsland is an associate professor in computer science and the postgraduate director in the School of Engineering and Advanced Technology at Massey University in Palmerston North, New Zealand. He has a degree in mathematics from Oxford University (1998) and a PhD from Manchester University (2002). Before moving to Massey University in 2004, he held postdoc positions in the UK, the USA, and Germany. His research interests are in mathematical computing, principally shape spaces, Euler equations, machine learning, and algorithms.