

The Massey University Smart Environments (MUSE) Manifesto

<http://muse.massey.ac.nz>

October 2010

1 Aims

While the term ‘smart home’ can mean many things, we have focussed our attention on a house that can monitor its inhabitants through unobtrusive sensors (such as motion detectors, cupboard door sensors, and electrical usage) and identify the particular behaviour that they are currently involved in. Note that we are focusing on unobtrusive sensors: while cameras and microphones can provide a lot of information about a person’s behaviours, they raise serious privacy and security concerns.

Once the behaviour has been identified, it can use this knowledge for a number of purposes:

Assurance Making sure the individual is safe (e.g., does not leave on the stove unintentionally) and performing routine daily activities (e.g., follows basic hygiene routines).

Support Helping individuals compensate for cognitive impairment (e.g., reminding them to take their medication if they have forgotten to do so).

Assessment Determining the physical or cognitive status of the individual (e.g., recognising if an abnormal behaviour is performed).

The last of these is of particular interest, since identifying things that are abnormal in a behaviour could highlight disease or uncertainty, and potential dangers that should be brought to the attention of a carer. We believe that it is important that the house *learns* about the behaviours of its inhabitant. Initially it will have very little knowledge about human behaviours – potentially even none – and then learns about behaviours as the inhabitant engages in them. This is a research problem in machine learning, and the main focus of our current work, as is described below.

In addition to the sensors that detect the actions of the inhabitant, the house can also be equipped with ambient sensors measuring temperature, humidity, etc. and it will also know the time that events occurred and the locations, since these are implicitly provided by the sensors. It is our belief that combining this contextual information with the actual sensor patterns should enable the detection of abnormal behaviours that could signify confusion or potential danger, without too many false positives (the ‘crying wolf’ problem).

2 Current Activities

We currently have ongoing research on several fronts:

Defining the precise requirements of a smart home We have identified a collection of *use cases* that specify particular desired behaviours (and potential problematic responses) of the home. These are a work-in-progress that we are hoping that others will contribute to, and are available at <http://muse.massey.ac.nz/shmuc>.

Behaviour recognition Our current main focus is on using the sensor information to identify the behaviours that the inhabitant is currently involved in. We have considered a variety of machine learning approaches to this problem and evaluated them on a publicly available smart home dataset to see how they deal with noise, variation in the ways that behaviours occur, and other issues of real-world datasets.

There are two principal approaches to the problem of behaviour recognition, based on whether the dataset is initially annotated by a human so that there are labels identifying the correct behaviours. While having the labels is a simpler and more reliable (albeit, still difficult) problem, it is also time-consuming for the inhabitant, and unlikely to be performed in a widely-used smart home system. We are therefore also investigating methods by which the system can learn without the requirement of any human generated labels for the data.

Reasoning about behaviours Once behaviours have been recognised, the home needs to analyse them to see whether or not they are correctly performed, and whether they are potentially dangerous. This requires a system that can reason about the behaviours and the effect that sets of behaviours undertaken at different times can have on each other.

Using contextual information There are some behaviours that are perfectly normal in themselves, but become abnormal because of when or where they are carried out: using the heater when the temperature is 30 degrees Celsius in summer, or eating in the bedroom, could be examples. By adding this information to the behaviour recognition system, we are finding additional ways to detect abnormality in the behaviour of the house inhabitant.

Giving feedback to the house In order for the house to learn what is and is not a reasonable output, it needs to be able to receive feedback from the user, whether this is the inhabitant or the carer. While this can be as simple as a console that people can log in to (for example, via the internet) or communication through SMS messages, we are also looking into emotion recognition, with the hope that the house will be able to identify what emotions its actions cause, and so reinforce its behaviour as required. (Note that this is likely to require a variation on our intended aim of avoiding cameras, as emotion detection without images is notoriously difficult. This project is thus slightly separate to our other research foci.) Additionally, emotion recognition can be a useful addition input to the system, such as for detecting when the inhabitant is unhappy.

Sensors for smart homes We are starting to consider what types of sensor can be used within a home, that can be unobtrusively installed into existing homes, and what information they can relay to the house computers.

3 Areas of Collaboration

Currently we do not have any ability to collect our own data, and have to rely upon a single publicly available dataset collected over two weeks in the house of a healthy inhabitant in America. This is obviously rather limiting. We want to be able to evaluate and test our algorithms in the New Zealand context with people in the correct sectors of society, and we cannot do this on our own.

Additionally, while we have a number of postgraduate research students currently working with us, there is always the potential for more where funding allows. If these students are also able to interact with people working in the domain of home-based care, then this will be additionally useful.

4 Limitations and Caveats

We are an academic research group, primarily interested in research into the computer science and artificial intelligence aspects of the problem. While we would love our research to be useful, we do not produce complete products, nor do we do know anything about the medical ethics or related topics for this area.

5 More Information

For further information, please visit our website <http://muse.massey.ac.nz>. You can also contact us by email (h.w.guesgen@massey.ac.nz or s.r.marsland@massey.ac.nz). We are happy to supply copies of academic publications on request.