Published in the Special Issue of Pervasive and Mobile Computing Journal on ``Human Behaviour in Ubiquitous Environments”, Elsevier B.V, Volume 6, Issue 4, 2010.

Title: Affect-aware behaviour modelling and control inside an Intelligent Environment

Author(s): Enrique Leon[1], Graham Clarke[2], Victor Callaghan[2], Faiyaz Doctor[2]

Affiliation(s):

[1] Fatronik

[2] Department of Computing and Electronic Systems, University of Essex

Full address for correspondence, including telephone and fax number and

E-mail address:

Department of Computing and Electronic Systems

University of Essex

Wivenhoe Park, Colchester, United Kingdom CO4 3SQ

Tel. No: +44 1 206 872682

Fax No: +44 1 206 872788

E-mail: graham@essex.ac.uk

Published in the Special Issue of Pervasive and Mobile Computing Journal on ``Human Behaviour in Ubiquitous Environments", Elsevier B.V, Volume 6, Issue 4, 2010.

Affect-aware behaviour modelling and control inside an Intelligent Environment

Enrique Leon, Graham Clarke, Victor Callaghan, Faiyaz Doctor

Abstract— From its inception research in Affective Computing has inspired a new generation of applications aimed at improving the interaction between humans and machines. Car driver concentration monitoring, Adaptive Computer Interfaces that react to affective states, e-learning systems that convey students' emotions and homes that cater for affective therapy are all examples of affective technology that could soon be part of our daily lives. In this paper an exploratory study is presented providing initial evidence of the positive effects of emotional information on the ability of intelligent agents to create better models of user preferences inside smart homes. An experiment was carried out over an 8-day period inside the iSpace which is a self-contained apartment used as a testbed for various software agents paradigms. Preliminary results suggest that an agent incorporating *valence-based* emotional data into its input array can model user behaviour in a more accurate way than agents using no emotion-based data or raw data based on physiological changes. Improved modelling of human activities is of great significance for pervasive computing as it leads to improved user satisfaction and more effective utilisation of resources.

Key Words: Affective Computing, Context Awareness, Intelligent Agents, Intelligent Inhabited Environments.

1. Introduction

The traces of emotions found in decisions sealing the fate of a country, the successful accomplishment of a scientific endeavour or the development everyday human relationships proved for centuries, at least intuitively, that emotions play a significant role in our behaviour. It was not until the nineteen nineties, sometimes known as "decade of the brain", that researchers from various disciplines provided strong evidence as to how emotions influence reasoning in our decision making as well as our motivational and learning mechanisms. For instance, neurologists demonstrated that emotions sometimes override reasoning in situations demanding quick decisions and immediate actions (Ledoux). They also evidenced that affective states are an important neurological regulator of the relationships between humans and their environment and that normal behaviour is greatly disturbed in the absence of such regulators were individuals show an inability to observe “socials conventions”, a tendency to take actions adverse to one’s own well-being that lead to financial or inter-personal losses, and repeated engagement in disadvantageous actions showing disregard from previous mistakes. Jocelyn Pixley argues that the intrinsic uncertainty associated with financial markets makes emotions such as confidence and trust an unavoidable element in corporative decision making (Pixley, 2002). Marcer provides examples of how emotions might be useful in formulating better explanations of rationality in the context of allegiance formation, justice, and strategic choice. He declares that while emotions might lead to mistakes so does cognition (Mercer, 2005). Moreover, emotions are so closely linked to decision-making that they can also be an explanation of why political choice often violates Bayesian updating leading to low-quality decisions (Redlawsk, 2002).

It has now become clear that emotional and cognitive processes are two interrelated, cooperative, inter-dependant constituents of our being, rather than separate, incompatible, independent elements. This combination of cognitive and emotional factors is an ingrained part of human behaviour and determines the course of our actions and development (Music, 2001). The aim of this paper is to explore this idea from the perspective of computing systems that aspire to learn user activities and accommodate the environment to provide comfort, safety and energy efficiency. Our argument is that, if emotions participate in our decisions and actions, pervasive systems which functioning depends on accurate modelling of user behaviour may in principle be able to better learn and reproduce user activities using a combination of contextual and affective information. Furthermore, we claim that pervasive systems that make inferences based on contextual information alone may be at disadvantage of those that consider the emotional elements that seem connected with our actions..

1.1. Affective Computing

The term affective computing was coined by Picard in the mid nineteen nineties (Picard, 1995) to describe computer methods that are related to, derived from or are deliberately designed to influence emotions. It encompasses two areas: *emotion synthesis* used to artificially imitate some of the physical or behavioural expressions associated with affective states, and *emotion analysis* which is often employed in decision making for interactive systems. Emotion synthesis is useful to develop ways to communicate with humans at a subjective level involving social participation, for example using robots. Emotion analysis could be used to monitor the emotional state of a subject, taking actions based on the type of individual feeling being experienced.

Numerous examples of the use of affective processing for various purposes exist. Some of these relate to the detection of emotional states by means of facial, speech or physiological features (Kim et al., 2002, Nasoz et al., 2003, Picard et al., 2001, Avent et al., 1994, Rosenblum et al., 1996, Nicholson et al., 2000, Moriyama & Ozawa, 1999, DeSilva & Ng, 2000, Yoshitomi et al., 2000). Others have instead developed computing systems that are capable of imitating, experiencing and/or reacting to people’s emotions (Garzon et al., 2002; Morishima, 2000, Cañamero & Fredslund, 2000, Aylett, 2004, Itoh et al., 2004, Suzuki et al., 1998, Mobahi & Ansari, 2003, Healey et al., 1998, Ark et al., 1999). Just a few have pointed at the issue of affective computing in the context of pervasive computing (Adelstein et al., 2005, Schultheis, 2007).

1.2. The problem of behaviour modelling inside Smart Home

The term intelligent inhabited environments (IIEs) or Smart Homes has been used to denote the physical embodiment of Ambient Intelligence (AmI), i.e., pervasive computing confined to habitations and buildings rather than to open spaces. Various research initiatives involving IIEs have been suggested in the past including the Intelligent Home (iHome) at Massachusetts University at Amherst, the Aware Home Research Initiative at Georgia Tech, the Neural Home at the University of Colorado, the PlaceLab developed as part of the MIT’s House_n project, and the HomeLab, an endeavour initiated by Philips. An important requirement in devising an IIE is to being able to build a model of users preferences based on continuous observation of their behaviour. The numerous intricacies of human behaviour means this requirement is not easily met and efforts have been directed at advancing the capacity of IIEs to faithfully represent user actions. Some notable approaches include the work of Eng et al. (2004), Barger et al. (2005), Heierman et al. (2004), Mozer et al.(1998), Amigoni et al. (2005), and Rutishauser et al. (2005). Notably, none of these studies has considered the instrumentation of affective computer systems as part of IIEs machinery.

1.3. Towards the Integration of Affective and Pervasive Computing

The idea of being able to exhibit emotions through electronic means has captured the imagination of many researchers in various areas of computing including inhabited intelligent environments and robotics. In fact, the creation of artificial entities capable of displaying affect and interacting with users at an affective level represents fertile ground not only for computer science but also for medicine and psychology. Furthermore, many other social and technological fields such as sociology, politics, finance, aeronautics, and the automotive industry could benefit from the detection, utilisation, and eventual imitation of human feelings.

In this context, the hypothesis being explored in this paper is that emotion detection would enhance an IIE preference model by devising a more detailed representation of the user's motivations and their relationship with the environment. More specifically, the work presented here will be an attempt to answer a number of basic questions associated with emotion detection and its application to pervasive computing, such as: - Is emotion detection useful i.e. does it improve the way pervasive systems model user behaviour inside IIEs? Can emotion detection enhance adaptability and increase the agent's ability to adjust the environment to reflect user's habitual behaviour and thereby increase their comfort as measured by their satisfaction with the agent's management of the iSpace (see 2.1 below)? Ancillary questions include how emotion detection information should be included in pervasive systems. Should it be based on the raw indication of physiological changes or on high-level (pre-processed) dimensions of emotion? Other questions related to this issue involve the emotional clasess being employed – positive, negative and neutral and whether these are sufficient to improve user comfort?

In addressing these questions we have developed a personalised physiology-based affective model for the control of pervasive computing environments, in our case a smart home called the iSpace. In this respect, our work could be seen as an extension to Picard's concept of affective computing where emotional models are seamlessly integrated with the environment with the intention of facilitating and improving human machine interactions. In our work, the affective agent learns, utilises and, in effect, replicates the relationships between affective states and the environment in a non intrusive manner using observations arising from continuous interaction with the user.

Note, that whilst the integration of affective and pervasive computing addressed in this paper has been discussed in the IIE community for some time, this is the first time, to our knowledge, that experiments have been carried out using an affective wearable with the intention of probing the value of emotions for adaptive systems employed in pervasive computing.

2. Background

A trial was carried out inside the iSpace, an experimental agent-controlled smart home, to assess the degree to which emotional data could contribute towards improving the modelling of the user's behaviour. An enhanced representation of the way the user interacts with the environment could lead to better agent adaptability, reduced need for user intervention (manual operation of devices and actuators), more efficient use of resources and, ultimately increased comfort.

2.1. The iDorm and iSpace Testbeds

The Intelligent Dormitory and the iSpace developed by the Intelligent Inhabited Environment Group (IIEG) at the University of Essex are state-of-the-art testbeds for researching agent paradigms and pervasive computing methodologies. The rationale underpinning these facilities is to provide researchers the opportunity to explore new ways of users interacting with technology inside domestic environments. Underlying the iDorm and iSpace is an extensive sensor and actuator network that both provide the key functionalities together with capturing user behaviour thorough monitoring user activities. The network and middleware infrastructure is based on Transmission Control Protocol/ Internet Protocol (TCP/IP) and Universal Plug and Play (UPnP). These environments not only sense the behaviour of the user but also create rule based user profiles that are used to pre-emptively orchestrate the operation of sets of networked devices to provide an intelligent behaviour (e.g., based on the users habitual behaviour the iSpace can learn how to control the heaters in the most economic way or, in a care setting, might identify hazardous situations).

2.2. Real-time Emotional Valence Detection

The first step towards the effective incorporation of affective and pervasive computing into IIEs lies in the accurate identification of the emotional state of the individual being analysed. In previous experiments, the combination of Autoassociative Neural Networks (AANNs) and sequential analysis namely the Sequential Probability Ratio Test (SPRT) (Fu, 1968), proved to be effective to detect changes in physiological signals associated with neutral, positive, and negative emotions (Leon, 2004a, 2007). This novel approach to detecting emotional valence is based on principles previously used to identify and compensate for sensor failures in nuclear plants and aircrafts and relies on the idea that positive and negative emotions could be identified by measuring the physiological differences they demonstrate against the neutral state.

An AANN is a Back-propagated neural network with a special architecture designed to learn the identity function, i.e., outputs always approximate inputs (Kramer, 1992). Thanks to their distinctive structure AANNs possess an intrinsic noise and gross error filtering capability related to the relationships they establish between input variables. In this context, data outliers caused by sensing error or data corruption are continuously eliminated at the output. In order to achieve this kind of robustness an AANN must be trained using exemplars that characterize the natural input-output behaviour of the phenomenon being modelled.

The SPRT on the other side is a statistical method used to dynamically detect changes in the state of a given process using sequential sampling offering a balance between an optimal number of samples without increasing the probability of misrecognition (Fu, 1968). The SPRT has been shown and it is still widely regarded to be an optimal classification technique to dynamically determine whether a given input pattern belongs to either of two categories (Cheng, 2002).

In our method, an AANN is trained to memorize the physiological concomitants of the neutral emotional state using information previously acquired from a number of individuals under controlled conditions resembling regular domestic activities and emotional changes. Five measures stemming from the autonomic system (heart rate (HR), skin resistance (SR), blood volume pressure (BVP), gradient of the skin resistance (GSR) and the speed of the combined changes in the all variables’ incoming data (Signals’ entropy, CS)) were used in the present study (Figure 1).

As a result of the training, changes in data patterns that result from non-neutral emotional states can be dynamically identified by continuously calculating data *residuals*, i.e. the numeric differences between the actual inputs and their AANN estimated values. In real-time operation, the mean of such residuals is very close to zero for data associated with the neutral state. When input sensor values drift because of physiological changes caused by a positive or negative emotional episode, the mean residual values significantly deviate from zero. Two SPRT modules are then utilized to statistically classify such variations into an emotional dimension using features extracted from training data linked to the targeted emotional dimensions (see Figure 2). In order to improve recognition rates signals used both in the classification of emotional status and valence are pre-selected by means of a calculation of the instantaneous Davies-Bouldin Index (DBI) (Davies & Bouldin, 1979) on residuals obtained during AANN training. Note, that while only heart rate (HR) and signal entropy (CS) are involved in emotional classification all input signals are needed in order for the AANN to establish inter-variable relationships associated with the emotional dimensions.

The main advantage of the SPRT is that it requires a minimal number of measurements before being able to reach a conclusion about the two hypotheses being evaluated (neutral Vs non-neutral and positive Vs negative in this case) (Fu, 1968, Tartakovsky, 2001). Based on the continuous, accumulative evaluation of incoming residual values, the SPRT calculation stops whenever the likelihood ratio increases beyond an upper limit or decreases beyond a lower threshold. Such boundaries are established using the solution spaces associated with to two targeted classes under analysis (neutral and non-neutral, positive and negative) and the Probability Distribution Function (PDF) of the variable being analysed. In this way the SPRT modules differentiates between transitory physiological drifts and an actual emotion making the entire system robust to perturbations caused by non-emotional sensor behaviour. The number of samples (time window) needed to identify an emotional change may vary but in previous studies this number was at around 9 samples in average (sampling rate 20Hz) (Leon, 2004a). See Appendix 1 for more information on the algorithms used in the calculation of the SPRT.

Published in the Special Issue of Pervasive and Mobile Computing Journal on ``Human Behaviour in Ubiquitous Environments", Elsevier B.V, Volume 6, Issue 4, 2010.

This technique to identify emotional changes has been shown suitable for user-independent operation and demonstrated robustness against bodily alterations produced by variable degrees of affect intensity and the effects of light to moderate physical exertion (Leon et al., 2005, 2007).

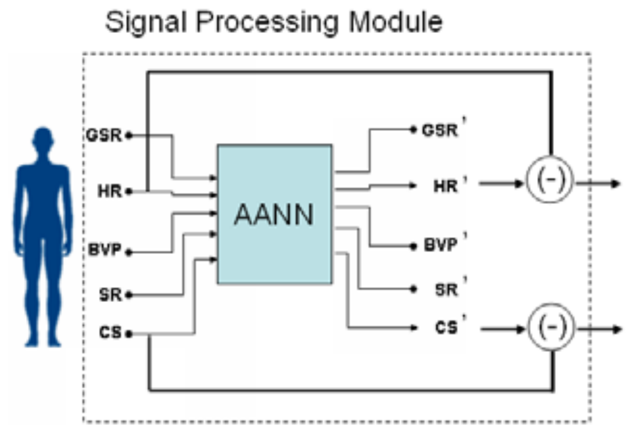

Figure 1. AANN operation in the Signal Processing Module.

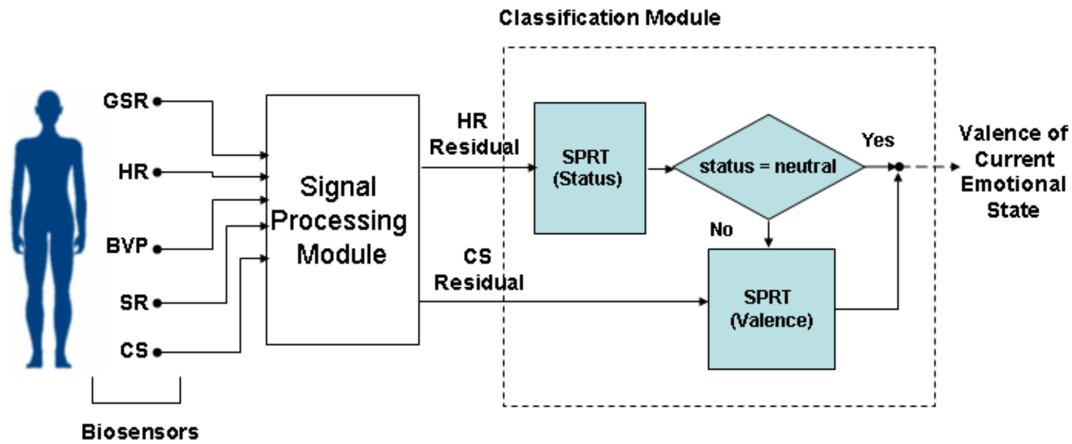

Figure 2. System architecture including the Classification Module.

2.3. eXperimental Vital sign-based Emotional State Transmitter (XVEST)

The XVEST is a wearable system capable of communicating the wearer's emotional state in real time using wireless technology and a finger clip with built-in sensors available from Discovogue Inc. Physiological data is collected via the finger clip and sent to a desktop computer using a Bluetooth connection where signals are processed to identify neutral, positive, or negative emotions using the pre-trained system described above. A software wrapper is used to present the emotion detection system as a UPnP device, making it compatible with other iSpace systems and allowing remote cross-platform access (see Figure 3). The XVEST mechanism makes possible reliable real-life experimentation.

Figure 3. The XVEST. Physiological signals are acquired via the fingerclip and sent to a PC using the Bluetooth transmitter (left). Batteries are inside lower pockets.

As a corollary to this overview of the XVEST and underlying methodologies we would like to mention that the Austrian psychologist Walter Mischel recognises two main techniques for the study of behaviour under real-life conditions: Verbal and non-verbal situational behaviour sampling and the measurement of physiological response concomitant to emotional reactions (Mischel, 1986). Thus, physiological emotion detection is not alien to the area of behaviour modelling.

2.4. Fuzzy Logic Agent

In this study, we used a fuzzy agent previously shown to deliver enhanced performance when operating inside intelligent environments (Doctor, 2005) against other agent architectures: the Adaptive On-line Fuzzy Inference System Agent (AOFIS). This is an intelligent mechanism that is not only able to generate and learn rules for the Fuzzy Logic Controller (FLC) but also the membership functions (the algorithms needed to “fuzzify” real-world data) associated with the environmental variables. After an initial *training phase* in which the user interacts naturally with the environment, AOFIS’s FLC learns a descriptive model of the user’s behaviour based on the event-related conditional rule data that has been accumulated.

Once a preliminary behavioural model has been constructed during the *training phase*, AOFIS initiates automatic pre-emptive control of the environment on behalf of the user (the *adaptation phase*) using the knowledge previously acquired. More specifically, using the combination of input/output pairs and the available membership functions, AOFIS creates an initial set of rules containing all the possible combinations of variables and membership functions. In a subsequent step, these rules are refined and compacted by the elimination of those parts of the rules that are associated with low membership values. During the *adaptation phase*, the agent continually monitors the environment and "fires" the rule that most faithfully reflects the current conditions and ambient state inside the iSpace.

Because AOFIS performance is based on observations over a limited period of the user's life (i.e. incomplete) the agent is designed to continually learn and adapt to new user actions. In order to avoid the incorporation of one-off activities, AOFIS behavioural model is not immediately altered a result of user actions. Instead, AOFIS adopts a principle called "learning inertia" in which rules are learnt only after they have re-occurred a number of times. Since each input/output pair can potentially result in one fuzzy rule created, the refined rule set is subjected to a process in which rules are merged together depending the similarities of their IF statement and also on the weighted value of their fuzzified components.

In our experiments, the agent's input vector comprised seven signals: the internal and external light levels, internal and external temperature, chair pressure, couch pressure and time measured as a continuous input on an hourly scale. Artefacts subjected to agent control included four variable intensity spot lights, a desk lamp, and two PC-based applications namely a word processing and a media playing program.. In the present study three different implementations of the AOFIS agent were compared: (1) the original agent with no emotional information - NON-E, (2) an agent using an extra input involving Fuzzy singleton discretised emotional values (1-Neutral, 2-Positive, 3-Negative) – DISCRETE-E, and 3) an agent with raw indication of emotional-based physiological changes added to the original input vector (three Gaussian Fuzzy sets stemming from the residual of the heart rate) - RAW-E. An example of a NON-E rule generated by AOFIS is (VLOW means very low):

```
IF
```

```
InternalLightLevel is LOW and ExternalLightLevel is HIGH and
InternalTemperature is VLOW and ExternalTemperature is LOW and ChairPressure is
OFF and BedPressure is OFF and Hour is Morning
THEN
ACTION_Light1_value is VLOW and ACTION_Light2_value is VLOW and
ACTION_DeskLight_state is OFF and ACTION_MSWord_state is STOPPED and
ACTION_MSMediaPlayer_state is STOPPED
```

Note that single actions reflect states and it is their accumulation over time which reflects behaviour. Several rules produced consecutively can be grouped to represent an action. In this paper we are not so much concerned with the internal algorithms governing agents operation, that is described elsewhere, rather we are interested in answering the more generic hypothesis of whether emotional data can improve the performance of agents managing pervasive computing systems, such as smart homes. Thus, for readers interested in understand the detailed workings of our environment control agent we refer them to our previous published work (Doctor, 2005).

It is worth mentioning at this point that the agent’s functioning is founded on two main principles: 1) “the user is the king” and can at any time effect changes on the ambient systems, and 2) the user should be oblivious to agent presence. In this context, realisation of agent/user interactions is twofold. Firstly by the agent’s automatic control of the user’s surroundings using knowledge about previously shown behaviour, e.g. switching the TV on automatically when the user is sitting in the couch; and secondly by the user operating the devices causing the agent to adapt its behavioural modelling, e.g. the user manually turning the TV off in the previous example.

3. Pre-Experimental Considerations

3.1. Continuous vs. Discretised

There are three main differences in terms of the way the Raw Fuzzified Emotional (RAW-E) and the Discretised Emotional (DISCRETE-E) agents operate: 1) The number of physiological inputs); 2) How emotional information is processed; and 3) How they respond to changes in physiological data.

The first difference lies in the number of bodily signals employed by the agents to gauge emotional valence. DISCRETE-E categorises emotions using residuals from two signals: Heart rate is used to identify a change from the neutral emotional status and valence is subsequently measured by means of the signal’s entropy (a characterisation of the signals temporal dynamics; i.e. how it changes over time). The reason for this two-phased operation is related to the implementation of the SPRT which permits only two hypotheses to be analysed (see Figure 2 above).

RAW-E, on the other hand, uses the output of the Signal Processing Module (HR residual) to measure changes in the user’s physiological state and relates such changes to the three emotional clasess in a single pass (Figure 1). For RAW-E such emotional clasess are given by the mean and standard deviation value of HR residuals calculated from three training datasets comprising data for neutral, positive and negative emotions. A DBI-based feature selection was also used on this occasion to identify HR as the attribute best distinguishing the three emotional classes. RAW-E’s use of parameters directly estimated from raw physiological data to classify emotions (the HR in particular) is akin to what researchers in psychology have done in the past, e.g. Ekman et al., (1983), Sinha et al. (1992). However, in contrast to these traditional approaches, the transformed solution space provided by the AANN offers enhanced inter-cluster spread of physiological data thereby improving class separation and further characterizing emotional classes (Leon et al. 2004b), a behaviour also preserved in HR residuals. It should be noted that although the agents’ internal operation, as well as the number of physiological inputs, differ from one another, the output for both cases provides a valence-based classification of the user’s affective state.

The second difference is founded on the membership functions employed to fuzzify physiological data. While RAW-E utilises Gaussian membership functions, DISCRETE-E employs singleton models which, in reality, represent rigid crisp sets. The reason for this lies in the type of emotional information DISCRETE-E and RAW-E can handle. On one side, the combined outputs of DISCRETE-E’s SPRT modules only permit four possible discrete values: 1 for neutral, 2 for positive, 3 for negative and 0 if a change that has been detected but no decision had been made at the moment of sampling. On the other side, RAW-E operates on the continuous non-linear value of the HR residual. As previously mentioned, the centres of RAW-E’s three Gaussian fuzzy sets are calculated based on the mean residual value of each of three emotional classes calculated from the data employed for AANN training.

Moreover, because Gaussian fuzzification permits different degrees of membership, RAW-E allows variable values of emotional valence to be measured (see Figure 4). For some emotions stronger body reactions might be equivalent to measuring varying affect intensity. For instance Dolf Zillmann has suggested that when the body is in a state of physical arousal, e.g. due to exercise or a lingering past emotional episode the expression of an emotion, such as anger or anxiety, is greatly intensified in the face of new triggering stimuli (Berkowitz, 1993).

The last difference is implicit to the way physiological data is processed. RAW-E has been designed to analyse the value of the "residual" calculated from the AANN real-time estimation of the heart rate. This value is a raw, immediate indication of how emotional-based physiological conditions are drifting from the neutral emotional state. It is immediate because as soon as the value of the residual of the HR enters into the solution space of the fuzzy membership function, it will indicate that an emotion has taken place. DISCRETE-E´s SPRT, on the other hand, employs a more stable mechanism that would only report a discrete emotional change after analysing the residuals of physiological signals from various sequential samples. In order words DISCRETE-E would not classify a transient sensor variation as an emotional change unless such variation lasts for a given period of time and correlates to an emotion in the combined analysis of all input signals. This is particularly useful for the study of noisy physiological inputs such as the BVP (especially under ambulatory conditions) since it is highly unlikely that non-emotional sensor behaviour will simultaneously produce an effect on all input signals that will be equivalent to the physiological response associated with an emotional change.

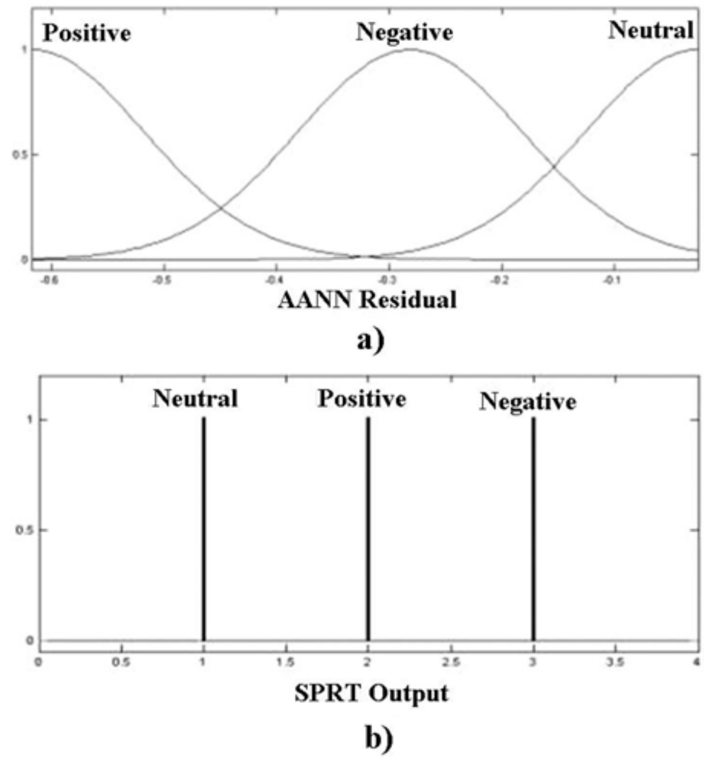

Figure 4. Membership functions for a) RAW-E and b) DISCRETE-E.

3.2. Intensive Small-sample Longitudinal Studies

There is an argument against longitudinal studies related to the fact that experimental subjects are often chosen in terms of their cooperativeness and availability which makes them unrepresentative and raises doubts about the generalisation of experimental results. However, “intensive small-sample” longitudinal studies, such as the one presented here, are sufficiently adequate in many areas including those associated with personal traits and physiological functions. The American academic Colin B. Hindley said about small-sample, intensive longitudinal studies:

> “The ultimate aim is naturally to produce generalizable findings, but reliance is placed not so much on having a strictly representative sample, which may be impossible to attain, as on the possibility of replicating important findings. This is the normal procedure in scientific inquiry, in which it is by no means assumed that any one investigation will produce conclusive matters”.

This study accords with such scientific aims. The experiment presented here will involve methods and produce initial evidence to test our hypothesis about the importance of emotional information in the adaptability of pervasive systems (as demonstrated by enhanced behavioural modelling). Consequently, the duration of our experiment and the number of subjects involved is directly related to the time needed to collect the numbers of samples required to show the effect we are investigating.

4. Method

A male participant aged 27 lived inside the iSpace for 8 days wearing the XVEST (see Figure 3). The XVEST’s fingerclip was firmly attached to the ring finger in the non-dominant hand by means of an elastic strap and, according to his account, did not impede the individual’s normal activities nor in any manner impaired his abilities to work on a computer. The first two days, the *training phase*, were used to collect ambience and emotional data to train the three fuzzy agents. Do not confuse this with training of the emotion detection system which is done offline. In the remaining 6 days, the *adaptation phase*, the participant performed a range of activities inside the iSpace comparable to those commonly undertaken in everyday life e.g., studying, eating, resting, exercising, etc. A crucial element in the present study is that the participant was asked to behave as naturally as possible and not to alter his normal behaviour or his response to unforeseen circumstances such as unexpected changes in the weather or his physical state e.g., in the event of feeling unwell. It is worth noting that the use of a single subject is well documented in affective computing where experimental design sometimes relies on several datasets produced from the same individual (Picard, 2001; Kurozumi et al., 1999; Ebine et al., 2000).

Figure 5. Experiments inside the iSpace.

In order for the comparison to be as accurate as possible, the three fuzzy agents were exposed to similar temperature and light conditions over the whole period of experimentation. Because of the restrictions imposed by the use of the actuators (they could only be operated by one agent at a time) parallel operation of the agents was not possible. Therefore, we decided that each agent would be used at a pseudo randomly selected time slot on the same day for the 6-day *adaptation phase* (*the learning phase* was the same for all the agents). In order to maximise data acquisition and not excessively restrain participant’s regular activities, the various time slots were chosen based on the times of the day in which more activity is likely to take place under normal conditions i.e., Morning or Breakfast time (8-10 AM), Midday or Lunch time (1-3 PM), and Evening or Dinner time (6-8:20 PM). The random time slot assignation was made with the condition that all the agents should be operated for the same amount of time. Thus, each agent was operated twice each morning, afternoon and evening yielding to sessions Morning1 and Morning2 of 7200s each, Afternoon1 and Afternoon2 of 7200s each, and Evening1 and Evening2 of 8400s each (Table 1). All sessions were accompanied by daily post-experimental debriefings in which the participant was interviewed and answered questions about the day’s experiences. Note that on the evening of the fourth day (DISCRETE-E’s Evening1) the participant felt unwell and remained lying on the couch for most of that session while DISCRETE-E kept controlling the room (performance results will highlight this issue when appropriate). Feeling sick is obviously a situation that might happen in any person’s regular living and thus its occurrence is useful for improving the agents behavioural model. In fact, because this is a situation considerably different from those found during training it would be interesting to see how DISCRETE-E coped with such abnormality.

		Agent Type		
Day/Time Slot		NON-E	DISCRETE-E	RAW-E
Day 1	8-10 AM	Morning 1		
	1-3 PM		Afternoon1	
	6-8:20 PM			Evening1
Day 2	8-10 AM			Morning1
	1-3 PM		Afternoon2	
	6-8:20 PM	Evening 1		
Day 3	8-10 AM		Morning1	
	1-3 PM	Afternoon1		
	6-8:20 PM			Evening2
Day 4	8-10 AM			Morning2
	1-3 PM	Afternoon2		
	6-8:20 PM		Evening1	
Day 5	8-10 AM		Morning2	
	1-3 PM			Afternoon1
	6-8:20 PM	Evening2		

Day 6	8-10 AM	Morning2		
	1-3 PM			Afternoon2
	6-8:20 PM		Evening2	

Table 1. Assignation of experimental time slots.

5. Results

5.1. Emotional Activity during Training

The routine undertaken by the subject inside the iSpace during agent training remained very much the same for the entire duration of the experiment. Daily routine normally commenced around 8.00 AM with the participant reading and responding to e-mail messages for a maximum of an hour and then moving on to perform some work-related activities including typing in documents and reading papers off the screen. This was done until the end of the morning session (10 AM) and continued for the totality of the afternoon session (1-3PM). The evening sessions usually had the user involved in leisure activities such as watching TV and playing videogames. According to DISCRETE-E results, as the day progressed and moved from less to more highly emotional charged activities (in accordance with the user's own account of the emotional content associated with each different activity, e.g., playing videogames was more emotionally charged than reading emails) the number of emotional changes (neutral-to-positive and positive-to-neutral) detected by the system also increased (see Table 2).

Time Slot	Emotional Changes (neutral-to-positive and positive-to-neutral)	
	Day 1	Day 2
Morning	18	0
Afternoon	32	4
Evening	45	34

Table 2. Total number of emotional changes during the *training phase* as detected by DISCRETE-E.

As mentioned earlier the initial emotional mapping is important since this will be used to construct the preliminary rule base subsequently refined and adjusted in the adaptation phase. The emotional information employed in producing such preliminary rule set is shown in Figure 6.

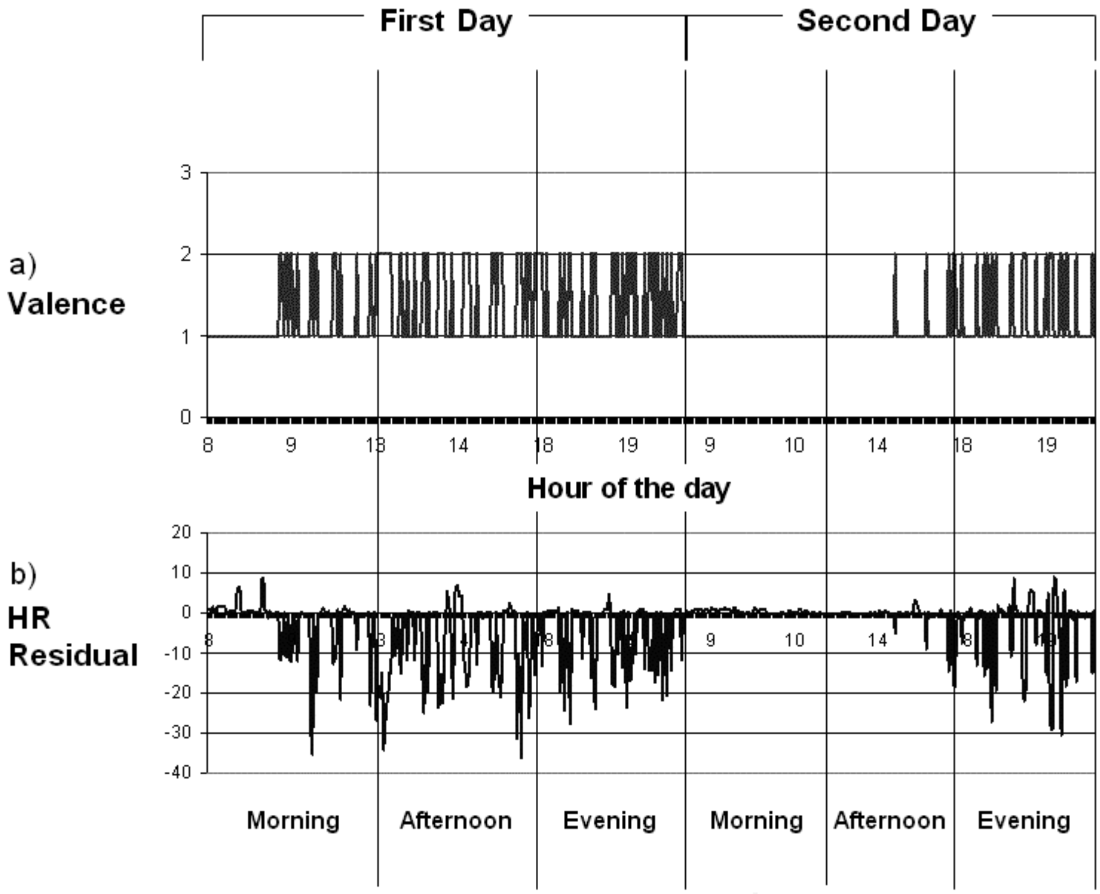

Figure 6. Emotional response during the two-day training phase expressed as a) DISCRETE-E's emotional classes and b) RAW-E's physiological data (AANN's HR residual).

Figure 6a illustrates the behaviour of the SPRT-based valence values on which DISCRETE-E's functioning was based. It is notable that for DISCRETE-E emotions remained in the neutral-positive valence values. This was corroborated by participant's self-reports which demonstrated that negative emotions did not occurred at anytime during the entire duration of the experiment. This is also in line with studies that reported that measured negative affect is infrequent and mild (Kahneman, 2004). Figure 6b on the other hand depicts HR residuals used in RAW-E's fuzzification as they varied across training sessions. It is evident that this measure suffers from sudden, large variations that seem to be the result of sensor disturbances (probably arising from body movement and/or sensor displacement) rather than an indication of emotional changes (it is difficult to imagine an emotional shift causing an abrupt drop of 35 bpm in HR). This does not mean that all residuals stem from noise but it is clear that RAW-E´s emotional input is a lot more erratic and unpredictable than the one used by DISCRETE-E. Nonetheless, there appear to be similarities in the operation of the two emotional agents as indicated by a high negative correlation (-0.876) between residuals and valences.

5.2. Initial Rule Base produced after *Training*

It was found that the way rules were constructed differed between the emotional and non-emotional agent. This difference indicated that emotional detection altered the way agents perceived actions inside the iSpace. Table 3 shows a group of rules created by the NON-E and DISCRETE-E from information collected around the same time in the morning of Day 1. It is worth clarifying that rules are produced sequentially in relation to an event that is taking place inside the iSpace. As such, the time passed between the creation of one rule and the next could take seconds, minutes or even hours depending on the user actions. That is, the user may in theory remain seated for hours without affecting any of the devices in his immediate surrounding or changes occurring in the environment. It is until a change in this scenario occurs that the agent will generate a new rule.

		NON-E		DISCRETE-E			
		Rule A	Rule B	Rule A	Rule B	Rule C	Rule D
IF	InternalLightLevel	VHIGH	LOW	VHIGH	MEDIUM	LOW	LOW
	ExternalLightLevel	VHIGH	VHIGH	VHIGH	VHIGH	VHIGH	VHIGH
	InternalTemperature	LOW	LOW	LOW	LOW	LOW	LOW
	ExternalTemperature	**MEDIUM**	**VLOW**	**MEDIUM**	**LOW**	**LOW**	**VLOW**
	ChairPressure	**ON**	**OFF**	**ON**	**ON**	**ON**	**OFF**
	CouchPressure	OFF	OFF	OFF	OFF	OFF	OFF
	Hour	MORN	MORN	MORN	MORN	MORN	MORN
	Emotion_Discrete	-	-	**POSITIVE**	**NEUTRAL**	**NEUTRAL**	**NEUTRAL**
THEN							
	Light1_value	VLOW	HIGH	VLOW	VLOW	VLOW	HIGH
	Light2_value	VLOW	VLOW	VLOW	VLOW	VLOW	VLOW
	DeskLight_state	**OFF**	**ON**	**OFF**	**OFF**	**ON**	**ON**
	MSWord_state	**RUNNING**	**RUNNING**	**RUNNING**	**STOPPED**	**RUNNING**	**RUNNING**
	MSMediaPlayer_state	**STOPPED**	**STOPPED**	**STOPPED**	**RUNNING**	**STOPPED**	**STOPPED**

Table 3. Comparison of rules created by NON-E and DISCRETE-E in the Morning of Day 1.

Considering rule A in both the NON-E and DISCRETE-E cases, it is apparent that environmental conditions (the rule antecedents) and also the settings of the actuators (the rule consequents) are exactly the same, the only difference being the extra emotional input in the case of DISCRETE-E. However, the rules that follow differ radically as a result of the DISCRETE-E picking up the emotional changes that occurred that morning. In DISCRETE-E’s rule B it can be observed that a change in the emotional state of the participant was also accompanied by a specific action (activation of the Media Player and deactivation of the MSWord application). This particular action that appears to be uninfluenced by other elements could not have been discovered by NON-E as it was diagnosed only because of its correlation with an alteration of the subject’s emotional state. NON-E’s Rule B seems to be associated with a later state where the internal light level is moved from VeryHigh (VHIGH) to LOW. DISCRETE-E’s rules C and D also seem to indicate the gradual change in lighting conditions.

As mentioned earlier, agent rules are generated based on the user’s behaviour as conveyed by his previous actions and will be used to assist the user whenever similar situations arise. The fact that the chair pressure sensor is off while MSWord is switched back on (see Rule D) does not imply that the agent is operating autonomously in an erroneous way but rather that the agent is trying to respond to current events taking place inside the iSpace. Note that the apparent inconsistency of MSWord re-starting while the user is separated from the chair is accompanied by another event: the activation of the desk lamp. Thus there appears to be a relationship between working on the computer and using the desk lamp that might be complementary to the chair/MSWord. This behavioural clue is, however, an act separate from the aforementioned emotional conduct.

We now consider another example with rules that were created in the same time frame during the evening sessions of Day 1.

		NON-E		DISCRETE-E		
		Rule Ab	Rule Bb	Rule Ab	Rule Bb	Rule Cb
IF	InternalLightLevel	VLOW	VLOW	VLOW	VLOW	VLOW
	ExternalLightLevel	VLOW	VLOW	VLOW	VLOW	VLOW
	InternalTemperature	LOW	LOW	LOW	LOW	LOW
	ExternalTemperature	MEDIUM	MEDIUM	MEDIUM	MEDIUM	MEDIUM
	ChairPressure	OFF	OFF	OFF	OFF	OFF
	CouchPressure	**ON**	**OFF**	**ON**	**ON**	**OFF**
	Hour	NIGHT	NIGHT	NIGHT	NIGHT	NIGHT
	Emotion_Discrete	-	-	**POSITIVE**	**NEUTRAL**	**NEUTRAL**
THEN						
	Light1_value	VHIGH	VHIGH	VHIGH	VHIGH	VHIGH
	Light2_value	VLOW	VLOW	VLOW	VLOW	VLOW
	DeskLight_state	OFF	OFF	OFF	OFF	OFF
	MSWord_state	STOPPED	STOPPED	STOPPED	STOPPED	STOPPED
	MSMediaPlayer_state	STOPPED	STOPPED	STOPPED	STOPPED	STOPPED

Table 4. Comparison of rules created by NON-E and DISCRETE-E in the Evening of Day 1.

Table 4 shows that NON-E's rules Ab and Bb and DISCRETE-E's rules Ab and Cb are almost identical (apart of course from the DISCRETE-E's extra input). However in between rules Ab and Cb, DISCRETE-E created another rule that reflected an emotional change. This emotional alteration (from positive to neutral) embedded in rule Bb seemed to have preceded the change in the couch pressure sensor that was later detected by the two agents. This extra rule created as a consequence of an emotional change provides additional information about the user behaviour and could therefore be useful to improve the identification and modelling of users typical actions. Although it is reasonable to assume that any additional variable offered to the agent's input array would provide finer discrimination of user's behaviour, the subtle behavioural clue represented in DISCRETE-E rules could have only been found by the inclusion of the emotional information, something that other type of sensors cannot provide. The potential overfitting effect of this extra variable in both RAW-E and DISCRETE-E´s input vector will be discussed in Section 6.

5.3. Performance Analysis

Because the agents were trained during two days only and, thus, were exposed to a limited amount of activities, a number of changes to the initial rule base are expected. However, since user behaviour remained relatively static over the 8 experimentation days, we expected to see rules being created rather than being changed, i.e. those activities already modelled by the agent should not be subjected to many modifications, should the agent have performed as anticipated. With this rationale in mind, agents' performance was evaluated according to two key categories: Interaction Model and User Comfort. These two parameters are related to the agent's direct interaction with the user and provide a clear indication of how well a particular agent succeeded in accommodating and/or adapting to user behaviour

5.3.1. Interaction Model

The interaction model is a depiction of how effective an agent was in modelling user activities inside the iSpace after the two training days and how well it adjusted the environment to the user's preferences during the 6-day adaptation period. This is gauged by calculating the number of training-generated rules that were modified, the number of rules that needed to be created and adapted during the *adaptation phase*, and the total number of online adaptations required over time.

Learning inertia means that rules are altered only if a situation the agent had previously captured in a given way demonstrates reoccurring changes. In this context, rule adaptations are made in response to two distinct types of events influencing only specific parts of the descriptive model learnt by the agent: 1) an apparent change in user acts that does not involve direct device manipulation, or 2) the user manually overriding agent control actions by operating devices such as lights or computer applications. Under the first scenario rule adjustments are the consequence of inconsistencies between the behaviour model the agent has previously created and current user actions. Such inconsistencies could be attributed to the agent's failure to represent some events faithfully or gradual, slight changes in user conduct. On the other hand, rule adaptations that derive from the user manually intervening to change the ambience are paired with poor agent decisions that result in user discomfort. The reason for this is that a rule altered this way indicates that the user has reiteratively requested changes in the environment that the agent has failed to carry out in a proper manner. At the same time, because of the aforementioned learning inertia effect, user interventions done to rectify agent operation will not cause an immediate change on the agents' rule base unless such interventions repeat a number of times. This condition would consequently delay an appropriate response from the agent and further exacerbate user dissatisfaction, an undesirable side effect of AOFIS learning. So, even when adaptations are intrinsic to life-long learning, too many changes to the rule base especially in a short period of time are a symptom of ineffective behaviour modelling.

In contrast, numerous unused rules are not detrimental for user/agent interaction but they do not possess useful information about lasting behaviour and as such they will eventually be discarded by AOFIS learning inertia. They might nonetheless have an impact on the agent's computational performance since they uselessly enlarge the rule base thereby hindering dynamic operation. So, in terms of the interaction model, better agent performance and behaviour modelling is represented by rules that are created, used and not subjected to frequent adaptations.

5.3.1.1 Adaptation of original rules

The number of rules produced during training that later suffered modifications relates to the efficacy of the agents to construct a rule base that would best represent user actions over limited period of time. Because the three agents used the same training data, the greater the number of adaptations to this initial rule base the poorer the preliminary modelling. Table 5 shows that in terms of the suitability of the initial FLC model, the agent with raw fuzzified emotional data (RAW-E) displayed greater accuracy since only 5.6% of the initial rules were adapted during training. In contrast, 9.5% and 17.4% of the initial rules generated by DISCRETE-E and the non-emotional agent (NON-E) respectively were modified. Claims about the advantages of RAW-E's initial model are also supported by the greater number of initial rules that were fired: 40.3% against 39.1% of DISCRETE-E and 27.2% of NON-E.

Category	NON-E	DISCRETE-E	RAW-E
Total Number of rules	**539** [423]	**768** [588]	**1345** [949]
Number of* initial *rules (Training phase)	***143***	***189***	***211***
Number of **initial** *rules adapted during adaptation phase*	***25***	***18***	***12***
% of Initial	**17.4**	**9.5**	**5.6**
Number of **initial** *rules that fired*	***39***	***74***	***85***
% of Initial	**27.2**	**39.1**	**40.3**

Table 5. Total and initial (before adaptation) rule base and number of fired and adapted rules. Numbers in brackets show results after data from session Evening1 has been removed.

5.3.1.2 Adaptation of new rules

Different to adapted rules, new rules are created as the result of the agent's continuous response to unseen events involving user actions, ambience settings and environmental conditions for which no previous record exist. These newly created rules should complement the existing rule base and enhance agent operation. Adaptations done on a large number of them demonstrate the agents difficulties in reproducing user actions, a disadvantageous "create-modify" dynamic.

Category	NON-E	DISCRETE-E	RAW-E
Number of **new** *rules generated during adaptation phase*	***396*** *[280]*	***579*** *[399]*	***1134*** *[738]*
Number of **new** *rules adapted during adaptation phase*	**162** [136]	**20**[0]	**633** [237]
% of New	**40.9** [48.5]	**3.4** [0]	**55.8** [32.1]
Number of **new** *rules that fired*	***350*** *[238]*	***535*** *[361]*	***940*** *[576]*
% of New	**88.3** [85]	**92.4** [90.4]	**82.8** [78]

Table 6. Rule base used during the adaptation phase showing number of fired and adapted rules. Numbers in brackets show results after data from session Evening1 has been removed.

The quality of the rules generated during the 6-day *adaptation phase* seems to have clearly favoured DISCRETE-E since only 3.4% of these were altered as opposed to 40.9% for NON-E and 55.8 % for RAW-E (Table 6). The accuracy of the new rules was also superior for DISCRETE-E since 92.4% of these were actually utilised compared to 88.3% for NON-E and 82.8% for RAW-E. The higher number of fired rules suggests that DISCRETE-E was able to adequately identify the subtleties in the user’s behaviour and adjust the model accordingly. In general terms, the interaction model generated by DISCRETE-E seems to possess improved overall consistency and accuracy in comparison to the other two since only 4.9% of its rules needed any adaptation while 79.2% of them were fired against 34.6% and 72.1% of NON-E and 47.9% and 76.2% for RAW-E respectively. Note that if we discard data for Evening1 the advantages of DISCRETE-E seem more dramatic especially in relation to NON-E.

5.3.1.2 Adaptations over time

Because AOFIS performs life-long learning it is expected that after the initial generation of rules, the number of adaptation will progressively diminish as the agent’s operation continues over time. In so doing the agent generation of new rules reduces to a stable low indicating that the rule base is generally adequate for modelling the behaviour of the user. The number of online adaptations thus measures the success of the agents in capturing user behaviour and adapting to his/her preferences. Thus, in a well designed system, the number of adaptation should decrease over time.

Session	Number of Online Adaptations		
	NON-E	DISCRETE-E	RAW-E
Morning1	88	12	120
Morning2	0	0	0
Afternoon1	0	8	0
Afternoon2	0	0	108
Evening1	26	0	396
Evening2	48	0	9
Total	**162**	**20**	**633**

Table 7. Number of online adaptations per session.

As can be seen from Table 7, in terms of rule adaptations, the DISCRETE-E did not need to make any further modifications to its interaction model after the first session and created fewer rules than NON-E and RAW-E. The NON-E on the other hand, performed a much better modelling in the afternoon sessions with no adaptations made or new rules created in either of the two sessions. The DISCRETE-E was in general a better agent with the least adaptations and the only one which, as expected, did not make more adaptations from one session to the other (see NON-E’s Evening sessions and RAW-E’s afternoon sessions).

5.3.2. User Comfort

Enhanced user comfort is one of the most important objectives of researchers in the area of IIE agents. User comfort could be evaluated by analysing the number of times the user had to interact with the system in order to adjust the settings inside the iSpace. It has been mentioned that manual adjustments mean that the agent failed to configure the environment to what the user expected under comparable internal and external conditions. Experiments indicated the superiority of the DISCRETE-E with only 10 user interventions for the entire 6 sessions (RAW-E and NON-E were overridden 21 times, an increase of 110 % in both cases) (see Table 8). If we discard data from the evening in which the user feel ill, the differences are reduced but still significant (DISCRETE-E, 10, NON-E 18, RAW-E, 18). If we compare the agents in a session-by-session basis the DISCRETE-E performed better than the NON-E and the RAW-E in 4 out of 6 (or 3 out of 5) instances

Session	**Number of User Interventions**		
	NON-E	DISCRETE-E	RAW-E
Morning1	4	1	2
Morning2	1	1	5
Afternoon1	1	6	2
Afternoon2	3	2	5
Evening1	3	0	3
Evening2	9	0	4
Total	**21**	**10**	**21**

.Table 8. Number of user interventions on 6 days of experimentation.

5.3.3. Overall Performance

Results in Table 9 demonstrate a clear advantage of DISCRETE-E in the categories related to the efficiency and quality of the rules encompassed in the Fuzzy Controller (Interaction Model) and also in the way user needs were satisfied.

	Overall Performance		
Category	NON-E	DISCRETE-E	RAW-E
Interaction Model (% of Adapted rules from Total)	34.6	**4.9**	47.9
Online Adaptations	162	**20**	633
User Comfort (No. of user interventions)	21	**10**	21

Table 9. Category winners.

6. Discussion

The fewer number of times the user had to override agent’s decisions in the overall results, along with a diminished need for rule alteration, suggests that the DISCRETE-E was able to establish a better representation of how the user behaved inside the iSpace. The reason for this could be that, just as the user responds to changes in the environment, emotions prompt individuals to act according to stimuli stemming from the various activities undertaken inside an intelligent inhabited environment. The reactions to such stimuli are not easily recorded by a non-emotional agent since they depend on the changes that occur in the user’s own psychological and physical perception of the surroundings. Such reactions provide crucial emotional information that offers a better representation of the various components that motivate and influence user actions (see Tables 3 and 4 above). Moreover, the lack of interaction at an affective level inhibits the symbiotic relationship between a non-emotional agent and the user and neglects important information about why and when certain events usually occur.

In other words, actions and decisions rooted on people’s emotions appear to be more likely to be captured by an agent sensing emotional change than one which is not aware of emotional changes such as the NON-E. As we have explained earlier, a consequence of our rule formation mechanism is that it enables the DISCRETE-E to identify the moment when an emotional change takes place thereby revealing user actions that otherwise would have remained undetected (for example, the change in emotions that preceded the shutdown of the word processor application).

6.1. Emotional vs. Non-Emotional

In terms of user satisfaction, DISCRETE-E was much better than the NON-E with 11 fewer interventions by the user. This advantage alone signifies the success of the experiments in terms of finding an agent that improves on another one. One might suggest that the DISCRETE-E performed better than the NON-E because of its increased granularity thanks to a greater number of inputs which potentially provides a more detailed depiction of user activities. However, the other agent with the same granularity attributes as the DISCRETE-E - the RAW-E - performed equally poorly as the NON-E in terms of the need for user intervention. Thus a greater number of sensors means more granularity and more rules but that does not necessarily imply better user comfort.

To elaborate on this last point, a useful insight can be gained by temporally ignoring the fact that DISCRETE-E contained emotional information. In pervasive computing we are always looking for methods to enhance the way intelligent agents operate. In this regard, the selection of the inputs on which such agents base their decisions is crucial. This task however is not easy: Some sensors would indeed improve agent performance but some others could in fact have the opposite effect and disturb operation by introducing noise, increasing operation time, and complicating decision making. Thus, the fact that we were able to find a sensed input (that from DISCRETE-E´s) that improved the levels of user satisfaction provided by an already tried and tested agent is encouraging.

6.2. Continuous vs. Discretised

Based on the performance results shown in Tables 5-8, it is plausible to suggest that, if an increase in the number of input sensors appears to be associated with more rules, it does not seem to have a direct effect on whether the agent is capable of learning from the user. For example, if only the two emotional agents with the same number of input signals were compared, the superiority of the DISCRETE-E in user comfort and overall interaction model (rule adaptations) is still manifest thus indicating that not only the quantity but also the quality of the information determines the agent’s performance.

But, why did the RAW-E not perform as well as the DISCRETE-E if both agents had the extra emotional information with RAW-E’s physiological and the DISCRETE-E’s emotional changes having a strong correlation indicating similar behaviour (see Figure 4)? Based on the results from the initial model the effect of variability in AANN residual on RAW-E performance can be discarded as a viable reason. The rationale for this lies in the fact that the interpolation provided by the RAW-E’s Gaussian fuzzy sets inherently implies a filtering in the input signals since close sensor or residual values would produce similar rules. Furthermore, all the agents are designed to eliminate duplicated or inconsistent rules and thus many rules that were generated as a consequence of the noise are continuously discarded.

The answer to the above question however could lie in another element associated with AANN residuals, extra granularity (remember that Gaussian fuzzy sets would fuzzify different levels of physical arousal). RAW-E was the agent that created the greater number of rules (1345) of the three available models (see Table 5). RAW-E’s capacity for rule creation seemed to work to developing an enhanced initial behavioural model but did not perform as well during the *adaptation phase* since many rules did not seem to reflect durable changes with only 76% actually being utilised. Nonetheless the number of fired rules did not seem to add much to the comparison between the agents since differences are not statistically significant. However the extra granularity achieved by RAW-E’s through the use of

AANN residuals could have in reality been a destabilizing factor rather than making control more accurate. More specifically HR mean residual values located in the intrinsic overlapping of Gaussian fuzzification could have blurred the distinction between situations in which rules applicable to a given activity were similar but not equal in terms of emotional content. That is, two comparable actions could have been respectively associated with a positive and a negative emotion because residuals were in the overlapping zone. In fact, fuzzified residual values close to emotions' HR thresholds are difficult to differentiate in terms of whether they stem from motion, sensor corruption or an actual emotional state (see Figure 6b).

In other words, very close HR values could have been associated with inexistent or contrasting emotional classes while all remaining input parameters showed similar values. This ambiguity in the information about a particular activity could have led to erroneous RAW-E actions during the 6-day *adaptation phase* especially at the moment of deciding which rules to use. This phenomenon consequently promoted the creation of new rules to compensate for those that were ineffective. DISCRETE-E´s use of singleton fuzzy sets is a straightforward solution for this dilemma since neighbouring values in this kind of fuzzification is non-existent and thus leads to more specific, accurate rules.

Although RAW-E's contrasting performance in the *training* and *adaptation* phases might seem to be a typical case of data overfitting, various elements depicted in Tables 5 through 7 rule out this hypothesis. First of all, overfitting might be an explanation for RAW-E generating too many rules but not to its relative failure in the user comfort analysis. For instance, NON-E had fewer inputs and yet it struggled to respond to the user preferences. This means that despite the RAW-E having a larger number of rules (produced during the *training* and *adaptation* phases), it failed to create an accurate model of user actions and did not produce as good a behavioural model as the DISCRETE-E. Second, AOFIS agents are not fixed inference models. They perform continuous online adaptations by replacing, modifying or creating new rules that in theory would improve the initial behavioural model. Thus, we would expect a greater number of user interventions during RAW-E´s first sessions which would gradually reduce as the agent makes adjustments to the initial model and discards flawed rules. This however was not the case and RAW-E was overridden more often in the second sessions (see Table 8). This is also an indication that enhanced “resolution” or granularity did not succeed in enabling RAW-E´s to come up with a better modelling in the long term.

In summary, we could then develop two possible explanations for DISCRETE-E’s superiority. Based on our experimental results, it could be argued that the link between user actions and emotional states is better modelled by the agents using high-level emotional labels (positive, negative and neutral) rather than raw indications of physiological changes. That is, user actions are better described and interpreted in terms of accepted emotional labels than through variable, mixed measures of emotional changes derived from raw physical state.

On the other hand, the modest performance of RAW-E particularly in the category of user comfort reveals that the sole addition of sensing information, the user’s physiological state in this case, into the agent’s input vector does not guarantee improved modelling of user’s activities. It is the inclusion of *meaningful* emotional data that provides a valuable insight not only into the current activities but also into the relationship between ambience conditions and the user’s state of mind.

7. Conclusions

The experimental results suggest that the utilisation of emotional data can improve the performance of Smart Home agents, particularly in those categories involving the modelling of user activities and comfort. Emotions provide the agents with a more accurate model of why and when the user undertakes certain activities at certain times. It is apparent that by knowing the affective motivations that accompany specific actions it is possible to endow software agents with a more accurate representation of the events taking place inside an intelligent inhabited environment.

In summary, our research has provided the following answers to the questions posed at the beginning of this paper:

- “Is emotion detection useful i.e. does it improve the way pervasive systems model user behaviour inside IIEs?” Yes, it was found that the model created by DISCRETE-E using behavioural and emotional information did require a lot less modifications than the one from a non-emotional agent based on behavioural data only.

- "Can emotion detection enhance adaptability and increase the agent's capability to adjust the environment to reflect user's habitual behaviour and thereby increase their comfort?" Yes, there is evidence that during the time the DISCRETE-E was controlling the iSpace, fewer manual interventions (the user manually changing ambience settings using the available actuators) were required, thus indicating that the user was satisfied with the way the DISCRETE-E controlled the environment. Although it is possible to argue that DISCRETE-E was controlling the iSpace in the fourth night when the participant felt unwell (see Table 1) and this alone could have contributed to fewer user interventions on that session, it is evident that the DISCRETE-E showed a consistently better (or at least similar) performance than the other two agents in 4 out of the 5 remaining sessions. Using the same training data and operating under comparable circumstances, the DISCRETE-E needed fewer adaptations to the model and also fewer user interventions which indicate better learning. Was there any effect on the DISCRETE-E derived from the user feeling sick? In theory, because Evening-1 was the only outlier in the user's conduct, DISCRETE-E should have created new rules to compensate for the lack of knowledge about what was happening. Indeed on that session, the DISCRETE-E did not make any adaptation but created 180 new rules (compare to NON-E's 116 and RAW-E's 396).

- "Ancillary to these questions is the question of how emotion detection information should be included in pervasive systems. Should it be based on the raw indication of physiological changes or on high-level (pre-processed) dimensions of emotion?" Figure 4 is a clear indication that HR residuals estimated from raw physiological data are intrinsically noisy. The fuzzification of such residuals by means of Gaussian functions seemed to promote the creation of an excessive number of rules that had a direct (negative) effect on the agent's performance. This was inferred from the fact that the RAW-E agent needed a greater number of adaptations to its behavioural model (see Tables 5-7). It is important to note that it is not the difference in membership functions what makes RAW-E to perform poorly against DISCRETE-E but the nature of the emotional information they handle. While SPRT modules provide a smoother, more stable measure of physiological changes, the simple quantification of residual's statistics is highly affected by non-emotional sensor behaviour. In this respect the DISCRETE-E's high-level emotional classes provided by the AANN/SPRT mechanism worked better.

- ”Other questions related to this issue involve the emotional dimensions being detected – positive, negative and neutral and whether these are sufficient to improve user comfort?” For the particular experimental setup presented here it is manifest that, based on the measures described earlier, a crisp valence-based tripartite emotional classification seems adequate to produce a better model of the user activities and also greater comfort than a non-emotional agent and one that uses raw measures of physiological changes. We therefore suggest that the use of emotional states in pervasive environments should be founded on discrete emotional classes, e.g. negative, positive, anger, sadness, happiness, etc. The use of features extracted from raw physiological measures does not seem to be the appropriate for real-time emotional interaction inside Smarthomes.

The findings presented here provide encouraging evidence of the importance of emotions for the analysis of human behaviour and information processing. However, some weaknesses and limitations inherent to the novelty of the experiments presented here are also palpable.

For the purpose of generalization, it might be desirable to extend the number of participants to include individuals from different age groups and who undertake a more varied set of activities. For example, the distinct psychological stressors that the elderly face are contrasting to those of young adults and might render the agents’ emotional and adaptive response clearly insufficient. It would also be desirable to compare the DISCRETE-E and RAW-E against non-emotional agents featuring more sophisticated adaptive mechanisms and larger input vectors that cater for supplementary contextual and physiological information such as the user’s location, respiration rate and/or temperature; it might be possible that additional behavioural information improves the performance of non-emotional agents. Also, it remains to be determined whether a DISCRETE-E that classifies emotions into additional subgroups (happiness, anger, sadness, etc.) would outperform such better-equipped non-emotional agent.

In the same manner, the present study only provides a partial, incomplete and somehow unbalanced answer to a more fundamental question: Do emotions permanently underlie our decisions and prompt us to act? We adhere to the view that emotions are essential for decision making but the evidence included here raises some questions. In our work we have assumed that the actions reflected in Tables 3 and 4 were associated with an emotional change but it is conspicuous that such actions also involved a degree of body motion. So, were the actions reflected in the rules the result of an emotional episode, e.g. the user stood up from the couch because he needed to go back to work and that is why his emotional state went from positive to neutral, or was the alleged emotional episode just a sensor perturbation caused by the motion involved in the action itself? We have explained our reasons to believe that the SPRT modules are sufficiently reliable to conclude that an affective change indeed occurred in the particular situations represented in Tables 3 and 4. The fact that many other rules, reflecting actions with no emotional changes, were produced seems to be additional proof of this. However, further studies and perhaps novel methodologies are needed to confidently uncover events that are purely sensor/motion-related and have no emotional content and those that arise from sentient phenomena. In this context there seems to be no other option but to perform long-term studies where affective computing is complemented with pervasive systems so that contextual information assists in differentiating such events. Thus, whilst forming an encouraging initial step, we hope this paper will motivate other, and perhaps larger experiments, which we anticipate will support these initial findings.

In summary, the specific objective of this work has been to contribute to a better understanding of emotions in the context of designing pervasive computing agents and gain currency towards a closer link between affective and pervasive computing. Finally, for our future work we are planning to refine these techniques, and deploy them in other pervasive computing environments. For example, we are collaborating with Shanghai Jiao Tong University (SJTU) in a project that is investigating how these methods might be used to improve that performance of pervasive learners using the open e-Learning platform at SJTU Network Education College (Shen et al., 2007). We also intend to investigate the use of affective pervasive computing in the context of policymaking, work relations and assistive technology.

Appendix I

Implementation of the SPRT

Considering that the measured parameter is a continuous function A(t) that should be categorized according to two stochastic processes A1(t) and A2(t), both possessing a normal distribution with means $\mu 1$, $\mu 2$ and standard deviation $\sigma 2$, the calculation of the SPRT at stage x is:

(1)

and the decision boundaries are given by

If

then A(t)= A1(t)

If

then A(t)= A2(t)

Where and are related to the probability error of misclassifying A(t) into process A1 and A2 respectively and are given by

Where α and β are the desired confidence values to recognize A1(t) and A2(t) respectively. Alpha (α) and beta (β) are selected in such a way that the system will choose A1(t) with at least (1-α) probability and A2(t) will be selected with probability at least (1-β). Very small values of α and β increase confidence in the recognition results but would require more data samples before moving to any of the two solution spaces. In some cases the values of the confidence intervals could be obtained from experimental data previously acquired. For the purpose of this work presented here a standard significance value of 0.05 for both alpha (α) and beta (β) was chosen.

Acknowledgments

We are indebted to Robin Dowling and Malcolm Lear for their invaluable support in the implementation of the XVEST. We also express our gratitude to Dr. Francisco Sepulveda for his contribution to our work and to the IIEG staff at the University of Essex for their support during experimentation inside the iSpace. This work has been funded by the EU IST Disappearing Computer funding program which has provided support for the iSpace and the Mexican National Council for Science and Technology (CONACYT). We also wish to acknowledge the valuable, enriching contributions of the anonymous reviewers of previous versions of this paper.

References

Adelstein, F., S. K. S. Gupta, G. Richard III, and L. Schwiebert: 2005, ‘Fundamentals of Mobile and Pervasive Computing’. New York, NY, USA: McGraw-Hill Professional Engineering.

Amigoni, F., N. Gatti, C. Pinciroli, and M. Roveri: 2005, ‘What Planner for Ambient Intelligence Applications?’. *IEEE Transactions on Systems, Man and Cybernetics Part A: Systems and Humans* **35** (1), 7-21.

Ark, W. S., D. C. Dryer, and D. J. Lu: 1999, ‘The Emotion Mouse’. *Eighth International Conference on Human Computer Interaction*, Edinburgh, Scotland, pp. 818-823.

Avent, R. R., C. T. Ng, and J. A. Neal: 1994, ‘Machine Vision Recognition of Facial Affect Using Backpropagation Neural Networks’. *Sixteenth Annual International Conference of the IEEE on Engineering in Medicine and Biology Society*, Baltimore, MD, pp. 1364-1365.

Aylett, R. S.: 2004, ‘Agents and affect: why embodied agents need affective systems’. *Third Hellenic Conference on AI*, Samos, Greece, pp. 496-504.

Barger, T. S., D. E. Brown, and M. Alwan: 2005, ‘Health-Status Monitoring Through Analysis of Behavioral Patterns’. *IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans* **35**(1), 22-27.

Bechara, A.: 2004, ‘The role of emotion in decision-making: Evidence from neurological patients with orbitofrontal damage’. *Brain and Cognition* **55**, 30-40.

Berkowitz, L.: 1993, ‘Aggression: its causes, consequences, and control’. New York, NY, USA: McGraw-Hill.

Cañamero, L. D. and J. Fredslund: 2000, ‘How Does It Feel? Emotional Interaction with a Humanoid LEGO Robot’. In: K. Dautenhahn (eds.): *The Human in the Loop. AAAI 2000 Fall Symposium*. Menlo Park, CA: AAAI Press, pp. 23-28.

Cheng, Q., P. K. Varshney, K. G. Mehrotra, and C. K. Mohan: 2002, ‘Optimal bandwidth assignment for distributed sequential detection’. *Fifth International Conference on Information Fusion*, Annapolis, MD, USA, pp. 1550-1556.

Damasio, A.: 1995, ‘Descartes' Error: Emotion, Reason, and the Human Brain’. New York, USA: Avon Books.

Davies, D.J. and D. Bouldin: 1979, ‘A Cluster Separation Measure’. *IEEE Transactions in Pattern Recognition and Machine Intelligence* **PAMI-1** (2), 224-227.

De Silva, L. C. and P. C. Ng: 2000, ‘Bimodal Emotion Recognition’. *IEEE International Conference on Automatic Face and Gesture Recognition*, Grenoble, France, pp. 332-335.

Doctor, F., H. Hagras, V. Callaghan: 2005, ‘An Intelligent Fuzzy Agent Approach to Realising Ambient Intelligence in Intelligent Inhabited Environments’. *IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans* **35**(1), 55-65.

Ebine, H., Y. Shiga, M. Ikeda, and O. Nakamura: 2000, ‘The recognition of facial expressions with automatic detection of reference face’. *Canadian Conference on Electrical and Computer Engineering*, Halifax, Nova Scotia, Canada.

Ekman, P., R. W. Levenson, and W. V. Friesen: 1983, 'Autonomic nervous system activity distinguishes among emotions'. *Science* **221**, pp. 1208–1210.

Eng, K., R. J. Douglas, and P. F. M. J. Verschure: 2005, 'An interactive space that learns to influence human behaviour'. *IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans* **35**(1), pp. 66-77.

Fu, K. S.: 1968, 'Sequential methods in pattern recognition and machine learning'. New York, NY, USA: Academic Press.

Garzon, M., P. Ankaraju, E. Drumwright, and R. Kozma: 2002, 'Neurofuzzy Recognition and Generation of Facial Features in Talking Heads'. *IEEE International Conference on Fuzzy Systems*, Honolulu, USA, pp. 926-931.

Goleman, D.: 1995, 'Emotional Intelligence. Why it can matter more than IQ'. New York, NY, USA: Bantam Books.

Healey, J., R. Picard, and F. Dabek: 1998, 'A new Affect-Perceiving Interface and its Application to Personalized Music Selection'. *Workshop on Perceptual User Interfaces*, San Francisco, CA, USA.

Heierman, E. O., M. Youngblood, and D. J. Cook: 2004, 'Mining Temporal Sequences to Discover Interesting Patterns'. *Workshop on Mining Temporal and Sequential Data, Tenth ACM International Conference on Knowledge Discovery and Data Mining*, Seattle, WA, USA.

Itoh, K., H. Miwa, M. Matsumoto, M. Zecca, H. Takanobu, S. Roccella, M. C. Carrozza, P. Dario, and A. Takanishi: 2004, 'Various emotional expressions with emotion expression humanoid robot WE-4RII'. *First IEEE Technical* Exhibition Based Conference on Robotics and Automation, Tokyo, Japan, pp. 35-36.

Kahneman, D., Krueger, A. B., Schkade, D. A., Schwarz, N., & Stone, A. A. (2004). A survey method for characterizing daily life experience: The day reconstruction method. *Science*, 306, 1776-1780.

Kim, K. H., S. W. Bang, and S.R. Kim: 2002, ‘Development of person-independent emotion recognition systems based on multiple physiological signals’. *Second Joint EMBS/BMES Conference*, Houston, USA, pp. 50-51.

Kramer, M.A.: 1992, ‘Autoassociative Neural Networks’. Computers and Chemical Engineering **16**(4), 313-328.

Kurozumi, T., Y. Shinza, Y. Kenmochi, and K. Kotani: 1999, ‘Facial individuality and expression analysis by eigenface method based on class features or multiple discriminant analysis’. *International Conference on Image Processing*, Kobe, Japan.

Ledoux, J.: 1998, ‘The emotional brain’. London, UK: Weidenfeld & Nicholson, The Orion Publishing Group Ltd.

Leon, E., G. Clarke, V. Callaghan, and F. Sepulveda: 2004a, ‘Real-time detection of Emotional Changes for Inhabited Environments’. *Computers and Graphics* **5**(23), 635-642.

Leon, E., G. Clarke, V. Callaghan, and F. Sepulveda: 2007, ‘A User-independent Real-time Emotion Recognition System for Software Agents in Domestic Environments’. *Engineering Applications of Artificial Intelligence* **20**(3), 337-345.

Leon, E., G. Clarke, F. Sepulveda, and V. Callaghan: 2004b, ‘Optimised Attribute Selection for Physiological Emotion Detection’. *Twenty-sixth Annual International Conference IEEE EMBS*, San Francisco, USA, pp. 184-187.

Leon, E., G. Clarke, F. Sepulveda, and V. Callaghan: 2005, ‘Real-time Physiological Emotion Detection Mechanisms: Effects of Exercise and Affect Intensity’. *Twenty-seventh Annual International Conference IEEE EMBS*, Shanghai, China, pp. 4719 – 4722.

Mercer, J.: 2005, ‘Rationality and Psychology in International Politics’. *International Organization* **59**, 77-106.

Mischel, W.: 1986, ‘Introduction to Personality. Fourth Edition: A New Look’. New York, NY, USA: CBS College Publishing, Holt, Rinehart and Winston.

Mobahi, H. and S. Ansari: 2003, ‘Fuzzy Perception, Emotion and Expression for Interactive Robots’. *IEEE International Conference on Systems, Man and Cybernetics*, Washington DC, USA, pp. 3918-3923.

Morishima, S.: 2000, ‘Real-time face Analysis and Synthesis using Neural Networks’. *IEEE Workshop on Neural Networks for Signal Processing*, Sydney, New South Wales, Australia, pp. 13-21.

Moriyama, T. and S. Ozawa: 1999, ‘Emotion Recognition and Synthesis System on Speech’. *IEEE International Conference on Multimedia Computing and Systems*, Florence, Italy, pp. 840-844.

Mozer, M. C., R. H. Dodier, M. Anderson, L. Vidmar, R.F. Cruickshank III, and D. Miller: 1995, ‘The neural network house: An overview’. In: Lars Niklasson & Mikael Boden (eds.): *Current trends in connectionism*. Hillsdale, NJ, USA: Lawrence Erlbaum Associates, pp. 371-380.

Music G.: 2001, ‘Affect and Emotion’ *Ideas in Psychoanalysis*. Cambridge, UK: Icon Books Ltd.

Nasoz, F., C. L. Lisetti, K. Alvarez, and N. Finkelstein: 2003, ‘Emotion recognition from Physiological Signals for User Modelling of Affect’. *Third Workshop Affective and Attitude User Modelling*, Pittsburgh, USA.

Nicholson, J., K. Takahashi, and R. Nakatsu: 2000, ‘Emotion recognition in speech using Neural Networks’. *Neural Computing and Applications* **9**(4), 290-296.

Picard, R., E. Vyzaz, and J. Healey: 2001, ‘Toward Machine Emotional Intelligence: Analysis of Affective Physiological State’. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **23**(10), 1175-1191.

Picard, R.: 1995, 'Affective Computing'. Tech. Rep. No. 321, Media Laboratory, Perceptual Computing Section, M.I.T.

Pixley J.: 2002, 'Finance organizations, decisions and emotions'. *British Journal of Sociology* **53**(1), 41-65.

Redlawsk, D.P.: 2002, 'Hot Cognition or Cool Consideration: Testing the Effects of Motivated Reasoning on Political Decision Making'. *Journal of Politics* **64**(4), 1021-1044.

Rosenblum, M., Y. Yacoob, and L. S. Davis: 1996, 'Human Expression Recognition from Motion using Radial Basis Function Network Architecture'. *IEEE Transactions on Neural Networks* **7**(5), 1121-1138.

Rutishauser, U., J. Joller, and R. Douglas: 2005, 'Control and learning of ambience by an intelligent building'. *IEEE Transactions on Systems, Man and Cybernetics Part A: Systems and Humans* **35** (1), pp. 121-132.

Schultheis, H.: 2007, 'IPRA - Enhancing the sensing abilities of ambient intelligence'. *Computational Intelligence* **23**(4), 464 - 483.

Shen, L., E. Leon, R. Shen, and V. Callaghan: 2007, 'Exploratory Research on an Affective eLearning Model'. *International Symposium e-Learning and Emerging Digital Tech (LET07)*, London, UK.

Sinha, R., W. R. Lovallo, and O. A. Parsons: 1992, 'Cardiovascular differentiation of emotions'. *Psychosomatic Medicine* **54**(4), pp. 422-435.

Suzuki, K., A. Camurri, P. Ferrentino, and S. Hashimoto: 1998, 'Intelligent Agent System for Human-Robot Interaction through Artificial Emotion'. *IEEE International Conference on Systems, Man and Cybernetics,* San Diego, CA, USA, pp. 1055-1060.

Tartakovsky, A: 2001, ‘Quickest Detection of Targets in Multiple-Resolution-Element Systems: Sequential Detection Versus Non-Sequential Detection’. *Workshop on Estimation, Tracking, and Fusion: A Tribute to Yaakov Bar-Shalom*, Monterey, CA, USA, pp. 326-350.

Yoshitomi, Y., S-Ill Kim, T. Kawano, and T. Kitazoe: 2000, ‘Effect of sensor Fusion for Recognition of Emotional States using Voice, Face Image and Thermal Image of Face’. *IEEE International Workshop on Robot and Human Interactive Communication*, Osaka, Japan, pp. 178-183.

Authors' Vitae

Enrique Leon received the BSc in Electronics from UNITEC, Mexico, his MSc in Software Engineering from University of Calgary, Canada (2002) and the PhD degree in computer science from University of Essex, UK (2007). He has been Head of Department for the Mexican Ministry of Public Education, a Software Consultant to REM Technology Inc., Canada, and Postdoctoral Fellow for Trinity College Dublin. Dr. Leon’s main research interests include affective computing, pattern recognition, and pervasive computing and their use in the social sciences. He is currently a Senior Research Officer with Fatronik, a major research centre for the development of intelligent devices.

Graham Clarke received his BSc in Architecture (Hons) from UCL, UK, the MSc in Applications of Computing from CNAA, University of North London, UK, the MA in Psychoanalytic Studies from Tavistock Centre/U. of East London, UK and his PhD (Psychoanalytic Studies) from University of Essex. He is Co-founder of the Intelligent Inhabited Environments Group at the U. of Essex in 1996. His research interests include ubiquitous computing and intelligent buildings 1996- present and affective computing 2002-present. Dr. Clarke worked as Computer officer at City University and Anglia Ruskin University and retired from Computer Science Dept. U. of Essex after 19 years as Computer Officer in 2005. Dr. Clarke is currently a visiting fellow in the Department of Computing and Electronic Systems at the University of Essex.

Victor Callaghan is Professor of Computer Science, director of the Digital Lifestyles Centre and leader of the Inhabited Intelligent Environments Group at the Essex University. He received a BEng and PhD from the University of Sheffield, UK in 1978 and 1982 respectively. His main expertise concerns intelligent embedded-agents, affective computing and end-user programming, especially applied to the creation of digital homes, intelligent buildings and other pervasive computing environments. Professor Callaghan has attracted in excess of £2 million of research funding, published over 100 papers and was responsible for setting up the University's Brooker Laboratory, Robot Arena and Intelligent Space (iSpace) research facilities. He is a key member of the organisational team for major pervasive computing conferences (e.g. Pervasive Computing & Applications, Ubiquitous Intelligence & Computing, Intelligent Environments) and journals (e.g. Pervasive Computing, Communications, Ubiquitous Computing & Intelligence and Intelligent Buildings).

Faiyaz Doctor received the BSc degree in Computer Science and Artificial Intelligence from University of Birmingham, Birmingham, UK, in 1998 and the MSc degree in computer science and artificial intelligent agents and the PhD degree in computer science from University of Essex, UK, in 2002 and 2006 respectively. Dr Doctor's research interests include computational intelligence, fuzzy logic, ambient intelligence, pervasive computing, and intelligent buildings. He is also interested in embedded agents and intelligent machines. Dr. Doctor is currently a Senior Research Officer with the Department of Computing and Electronic Systems Science, University of Essex, UK.