In Engineering Applications of Artificial Intelligence, Volume 20, Issue 3, Pages 337-345, 2007

A User-independent Real-time Emotion Recognition System for Software Agents in Domestic Environments

Enrique Leon (Corresponding Author)
E-mail: eeleon@essex.ac.uk
Graham Clarke
E-mail: graham@essex.ac.uk
Victor Callaghan
E-mail: vic@essex.ac.uk
Francisco Sepulveda
E-mail: fsepulv@essex.ac.uk

Department of Computer Science,
University of Essex
Wivenhoe Park, Colchester CO4 3SQ,
United Kingdom.
Telephone: +44 1206 872790, Facsimile: +44 1206 872788

Abstract

The mystery surrounding emotions, how they work and how they affect our lives has not yet been unravelled. Scientists still debate the real nature of emotions, whether they are evolutionary, physiological or cognitive are just a few of the different approaches used to explain affective states. Regardless of the various emotional paradigms, neurologists have made progress in demonstrating that emotion is as, or more, important than reason in the process of making decisions and deciding actions. The significance of these findings should not be overlooked in a world that is increasingly reliant on computers to accommodate to user needs. In this paper, a novel approach for recognizing and classifying positive and negative emotional changes in real time using physiological signals is presented. Based on sequential analysis and autoassociative networks, the emotion detection system outlined here is potentially capable of operating on any individual regardless of their physical state and emotional intensity without requiring an arduous adaptation or pre-analysis phase. Results from applying this methodology on real-time data collected from a single subject demonstrated a recognition level of 71.4% which is comparable to the best results achieved by others through off-line analysis. It is suggested that the detection mechanism outlined in this paper has all the characteristics needed to perform emotion recognition in pervasive computing.

Keywords: emotion detection, pervasive computing, software agents

In Engineering Applications of Artificial Intelligence, Volume 20, Issue 3, Pages 337-345, 2007

A User-independent Real-time Emotion Recognition System for Software Agents in Domestic Environments

Enrique Leon, Graham Clarke, Victor Callaghan, Francisco Sepulveda

Department of Computer Science, University of Essex, Wivenhoe Park, Colchester CO4 3SQ, United Kingdom.

Abstract

The mystery surrounding emotions, how they work and how they affect our lives has not yet been unravelled. Scientists still debate the real nature of emotions, whether they are evolutionary, physiological or cognitive are just a few of the different approaches used to explain affective states. Regardless of the various emotional paradigms, neurologists have made progress in demonstrating that emotion is as, or more important than reason in the process of making decisions and deciding actions. The significance of these findings should not be overlooked in a world that is increasingly reliant on computers to accommodate to user needs. In this paper, a novel approach for recognizing and classifying positive and negative emotional changes in real time using physiological signals is presented. Based on sequential analysis and autoassociative networks, the emotion detection system outlined here is potentially capable of operating on any individual regardless of their physical state and emotional intensity without requiring an arduous adaptation or pre-analysis phase. Results from applying this methodology on real-time data collected from a single subject demonstrated a recognition level of 71.4% which comparable to the best results achieved by others through off-line analysis. It is suggested that the detection mechanism outlined in this paper has all the characteristics needed to perform emotion recognition in pervasive computing.

Keywords: emotion detection, pervasive computing, software agents

1. Introduction

The aim of this work is to present evidence that user-independent real-time emotion recognition is feasible. From consideration of previous work using off-line analysis, it is suggested that Artificial Neural Networks (ANNs) and sequential analysis could be employed in real-time analysis to guarantee accurate classification. Moreover, ANN generalization capabilities mean that the proposed system could operate on never previously seen data and still provide good performance.

For a long time emotions have been the focus of widespread studies from a variety of scientific areas including medicine and psychology. In 1990, Salovey and Mayer developed a broad framework known as emotional intelligence (EI) to describe how humans perceive and utilize their emotions (Salovey and Mayer, 1990). The initial purpose of EI theorists was to investigate the significance of emotions within the context of intelligence, paying special attention to adaptation and behaviour. However, health, personality, personal ambitions and success have also been analysed from an EI perspective.

The importance of EI in the mechanisms governing human conduct has been highlighted thanks to the work of a group of neurologists who have investigated the idea that, in terms of regulating our decisions and actions, reason is consistently less significant than emotions and, furthermore, that a lack of EI could impair the relationship between humans and their environment (Goleman, 1995). It is evident then that, since emotion recognition is one of the most important components of EI and it has a direct effect on our ability to make optimal decisions (along with the ability to utilize emotions to make decisions), any attempt by computer scientists to model human interaction should, at least in part, be founded on an accurate identification of affective states. It is suggested that by ignoring the emotional component intrinsic to human decision-making, software developers have been missing valuable information that could potentially lead to inadequate interactive models.

One discipline that could particularly benefit from emotion detection is that of pervasive computing. Pervasive or ubiquitous computing involves the integration of computers into the environment allowing the user to interact with them in a more natural way. It is argued that by allowing embedded computers to recognize and use emotional information, software agents of the type used in Intelligent Inhabited Environments (IIE) should be able to use this information to better adapt to what the user wants, increase the accuracy of decisions derived from what the user does, and facilitate mutual interaction. Ultimately actions taken by affective IIE agents could be comparable to intelligent human activity, i.e., would guarantee user’s comfort.

1.1 Related research

In the last two decades there has been a considerable amount of research into methods for recognizing emotions using various physical parameters. Facial gestures, vocal utterances, and more recently bodily measures stemming from the autonomic nervous system and the brain have been used to classify affective states. It is suggested however that, because physiological measures are more difficult to conceal or manipulate than facial expressions and vocal utterances, and potentially less intrusive to detect and measure, they are a more reliable representation of inner feelings and remain the most promising way for detecting emotions in computer science.

Hitherto, the identification and classification of emotional changes has achieved mixed results ranging from 60-95.5% detection accuracy for facial recognition (Avent et al., 1994; Rosenblum et al., 1994; Sun et al., 2004; Bartlett et al., 2003; Anderson and McOwan, 2003; De Silva and Hui, 2003) to 50-87.5% for speech recognition (Nicholson et al., 2000; Tsuyoshi and Shinji, 1999), and 72% in bimodal recognition (face and speech) (De Silva and Ng, 1999). In physiological emotion detection some of the best results have been achieved by Kim et al. (2002) with 61.2% correct classification for 4 emotions, Nasoz et al. (2003) 50-90% for 5 emotions and Picard et al. (2001) with 81% for 8 emotions. Some of the recognition techniques employed in the above approaches include neural networks (Avent et al., 1994; Rosenblum et al., 1994; Sun et al., 2004; Bartlett et al., 2003; Anderson and McOwan, 2003; De Silva and Hui, 2003; Nicholson et al., 2000) and advanced statistical mechanisms (Tsuyoshi and Shinji, 1999; De Silva and Ng, 1999; Kim et al., 2002; Nasoz et al., 2003; Picard et al., 2001).

Some applications using emotion recognition for practical purposes include the “Affective DJ” - implemented as an affective wearable (Healey et al., 1998), a driver stress detector (Healey, 2000), a frustration level estimator for poorly-interactive computer applications (Fernandez, 1997) and an emotion-based auto-adjustable computer environment (Ark et al., 1999). Emotion detection and estimation has also been used to recognize emotional content in voicemail messages (Inanoglu and Caneel,

2005) and provide immediate medical feedback in tele-home health care systems (Lisetti et al., 2003).

It is important to mention, that one of the main omissions made by researchers investigating emotional states is that they repeatedly forget to verify whether subjects actually experienced the targeted emotional states at the measured level of intensity during experimentation. In many experiments, researchers often assume that participants react to emotional stimuli in a similar manner neglecting personal interpretation that sometimes means intended emotional states never occur. The supposition that individual responses to emotional stimuli are similar for all the participants and that the simple presentation of such stimuli suffices to elicit emotional states could lead to the acquisition of flawed or biased experimental data. A way of correcting this problem is by the use of self-reports collected after a given stimulus has been provided or by selecting experimental subjects whose emotional responsiveness would make it more likely that targeted emotions actually took place.

Another important assumption often found in emotional experimentation and in particular in physiological emotion detection is that situations causing physical arousal that are not linked to emotional episodes, e.g. exercise, can be ignored during data analysis and pattern recognition. The introduction of noise or the potential upheaval of physiological signals due to unforeseen stresses like exercise could make emotion detection more difficult and thus skew the results.

Finally, it should be said that the results achieved in the physiological methods described thus far are based on off-line analysis that requires collection of substantial amounts of data in order to estimate statistical features. As a consequence, on-line dynamic operation is difficult and even in those cases in which real-time implementation is possible (Nasoz et al., 2003, Lisetti et al., 2003), little attention is paid to issues such as robustness and user independence.

2. Background

In (Leon et al., 2004a), the combination of Autoassociative Neural Networks (AANNs) and sequential analysis, namely the Sequential Probability Ratio Test (SPRT), proved to be effective in detecting changes in 4 physiological signals, more specifically the electromyogram obtained from the masseter, blood volume pressure, skin conductance and respiration rate associated with emotional states from a single individual. The recognition rate on that occasion was 100%. This methodology is based on the idea that the detection of emotional changes using physiological signals could be likened to a real-time sensor validation process (Sensor Failure Detection, Identification and Accommodation (SFDIA)) in which emotional states could be detected by estimating the amount of deviation they demonstrate with respect to a neutrally-emotional state. SFDIA is commonly used in industrial processes to determine the moment in which a given sensor becomes corrupted, i.e., provides readings others than the ones expected under normal operating conditions. Alterations in the autonomic system associated with emotional states are identified by providing a Sequential Probability Ratio Test (SPRT) module with the continuous calculation of the difference between an actual sensor value and its AANN estimated counterpart.

Because the AANN is trained to mimic the input behaviour of the subject's neutral state, the mean of the difference is very close to zero (with a standard deviation similar to that of the noise introduced by the sensing device) when the physiological state of the subject is normal. When a given sensor value drifts because of a change in the physiological status of the subject which has been provoked by an emotional episode, the mean value of the residual deviates from zero. The SPRT value is consequently altered and the likelihood ratio displaced to either of the two solution spaces (neutral or non-neutral). Despite the fact that detection could be made using the physiological measure with the better classification attributes only, the relationship of all the parameters is needed for projecting the targeted variable into the AANN estimation space.

The main appeal of the use of sensor validation techniques to recognize emotional states is that there is no need to amass a great number of data samples before a decision can be made. This real time responsiveness is an important characteristic required by interactive systems in order to facilitate the acquisition of information and improve the quality of decisions made by agents in pervasive computing and particularly in Intelligent Inhabited Environments (IIE).

In addition to immediate response, an IIE affective agent should prove to be robust to physiological changes stemming from physical activities rather than emotional episodes and also to the different levels of affect intensity across individuals. Previous studies demonstrated that the methodology described above was able to resist perturbations due to unforeseen bodily changes and also to various degrees of affect intensity (Leon et al., 2005a).

In this paper this model is extended in various ways: First, instead of using data from a single subject, emotional information used to train the AANN was acquired from several individuals. Second, the number of bodily signals involved in the analysis was also increased by one and the recognition module modified to recognize not only neutral and non-neutral emotional states but also positive and negative emotions. Finally, the real-time performance and robustness of this new and improved system was tested by presenting the AANN with never previously seen real-time physiological data.

The design of the overall system is shown in Figure 1.

[Figure 1 goes here]

2.1 Relevance to Emergent behaviour

The emotion detection mechanism discussed in this paper is not only a potentially valuable source of information in ubiquitous computing but is also an ideal candidate to be part of complex emergent systems. In the behaviour based approach to AI, the equivalent to reasoning and planning in traditional AI is produced by arranging for an agent to have a number of competing processes that are vying for control of the agent (Brooks 85) (Steels 91). The “sensory context” determines the degree to which any process influences the agent. Thus, as sensing is derived from what is effectively a non-deterministic world, the solutions from this process are equally non-deterministic and result in what is termed “Emergent Behaviour” (behaviours or solutions that emerged but were not explicitly programmed) (Callaghan 04). Thus, anything that affects the context can have a hand in this type of "emergent behaviour". In this respect emotions can be seen as another contextual influence that can be added to the pool of processes that influence the decision of an agent. In many respects this is not dissimilar to Minsky's "Society of Mind" thesis that might be interpreted as viewing the mind a set of competing processes vying for control (Minsky 85). From work on mobile robotics, behaviour based techniques, and their associated emergent behaviour, have been shown to be particularly well suited to dealing with environments that are difficult to describe mathematically. Such difficulties arise from the complexity of the variables or dynamics together with real-time physiological signals themselves that are subject to much uncertainty and noise. The behaviour based approach overcomes this problem by dispensing with the an explicit model of the world, rather using the world itself as its own model; rather more eloquently summed up by Brooks' as "the world is its own best model" (Brookes 85). Emotions fit this approach well as they are dynamic and difficult to describe in mathematical form. Perhaps, rather more speculatively, even our own experience of them suggests a somewhat turbulent interplay with vying feelings, emotions and logic! In these respects emotion processing has a natural fit with behaviour based systems and emergent behaviour.

3. Methods

3.1 Autoassociative Neural Networks

Autoassociative Neural Networks (AANNs) are a special type of back-propagation neural network (BPNN) designed with a specific architecture and trained to learn the identity function, i.e., outputs equal inputs (Kramer, 1992). AANNs are very robust to noisy or corrupted inputs because of their improved association mechanisms and they also posses improved filtering and generalization properties.

Because AANNs can provide estimations for both faulty and healthy inputs even in the presence of several abnormal or corrupted inputs, they have been successfully utilised to perform SFDIA (Lu et al., 2001) where failures are detected by calculating the error (residual henceforth) between each sensor value and its corresponding estimation (Hines et al., 1997).

3.2 Statistical Probability Ratio Test (SPRT)

In SFDIA the residual provided by the AANN is usually subjected to a statistical process in order to detect when a sensor is changing from its expected value. The SPRT has been shown to be an optimal classification technique to determine whether a given input pattern belongs to either of two categories, usually faulty or healthy. The main advantage of the SPRT is that is it requires a *minimal* number of measurements before being able to reach a conclusion about the two hypotheses being evaluated (Fu, 1968).

The SPRT is continuously estimated using the residual value and stopped when the value of the likelihood ratio reaches one of two predetermined mutually exclusive thresholds. Such boundaries are established based on the solution spaces related to two targeted classes and the Probability Distribution Function (PDF) of each variable. Considering that the measured parameter is a continuous function A(t) that should be categorized according to two stochastic processes $A_1(t)$ and $A_2(t)$, both possessing a normal distribution with means μ_1, μ_2 and standard deviation σ^2, the calculation of the SPRT at stage x is:

$$LOG(SPRT)_x = \frac{(\mu_1 - \mu_2)}{\sigma^2} \sum_{x=1} \left[A(t)_x - \frac{1}{2}(\mu_1 + \mu_2) \right] \qquad (1)$$

and the decision boundaries are given by

If

$$LOG(SPRT) > LOG(e^{A_1}) \quad \text{then} \quad A(t)= A_1(t)$$

If

$$LOG(SPRT) < LOG(e^{A_2}) \quad \text{then} \quad A(t)= A_2(t)$$

Where e^{A_1} and e^{A_2} are related to the probability error of misclassifying A(t) into process A_1 and A_2 respectively and are given by

$$e^{A_1} = (1-\alpha)/\beta$$

$$e^{A_2} = \alpha/(1-\beta)$$

Where α and β are the desired confidence values to recognize $A_1(t)$ and $A_2(t)$ respectively. Alpha (α) and beta (β) are selected in such a way that the system will choose $A_1(t)$ with at least (1-α) probability and $A_2(t)$ 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 especially in the case of historic errors derived from false alarms and broken sensors.

3.3 Cluster analysis

A class separation analysis based on the Davies-Bouldin Index (DBI) introduced above provides an insight into the amount of inter- and intra-cluster separation within a given data set and indicates the attribute(s) that contribute to an optimal separation of two or more classes (Davies and Bouldin, 1979). The DBI has been successfully utilised in studies that involve pattern recognition of physiological signals, where lower DBI indexes reflect a better class separation (Sepulveda et al., 2004).

4. Experimental Procedure

In the present work, the SPRT and AANN residual values were used to detect changes from a neutral to a non-neutral state and, once a change was detected, determine whether it was negative or positive. To do so, a database of physiological information related to neutral, positive, and negative emotional states was built and an AANN trained with the neutral information data contained in such database. It is suggested, that the trained AANN would project the three classes in different areas of the estimation model and the SPRT would detect the drifts in the residual value. Previous studies have demonstrated that the utilisation of AANN is a valuable and effective mechanism to increase inter-cluster separation related to emotional polarity (positive or negative) (Leon et al., 2004b) enhancing the possibility of successful recognition.

4.1 Data set description

Physiological data included information from 3 emotional states (Neutral, Positive, and Negative) collected from 8 individuals (5 women and 3 men) aged 26-48 and grouped according to whether they were acquired before or after physical activity. This tripartite emotional classification is rooted in the valence dimension of the three dimensional view recognized by theorists in emotional assessment (the other two dimensions being arousal and dominance) (Lang et al., 2001). This dimensional interpretation of emotions stems from factor analytic studies on semantic differential demonstrating that the variance in emotional assessment could be explained using three main dimensions

(Osgood et al., 1957). This simple standard categorization of emotions facilitates analysis and provides objectivity to this study avoiding the use of subjective labels that have given rise to much discussion among psychologists (depending on the emotional paradigm employed there are between 2 and 12 human emotions).

The proportion of data samples was approximately ¾ before exercise and ¼ after exercise. Based on their responses to the Affect Intensity Measure (AIM) questionnaire (Larsen et al., 1986) four of the subjects were considered to have high affect intensity (scored more than 0.5 standard deviations above the mean AIM value (Prkachin et al., 1999)) and the remaining four (3 men and 1 female aged between 26 and 35) were on the medium and low affect intensity scale. These physiological data have been employed in studies involving an analysis on the effect of affect intensity and physical exertion in the performance of the AANN. Based on the calculation of the DBI index and the Wilcoxon test, it was demonstrated that neither physical activity nor affect intensity play an important role in the associative and generalization properties of the AANN trained with neutral data (Leon et al., 2005b). Thus, it can be then inferred that the use bodily information from individuals with variable affect intensity collected under resting conditions or after physical activity should not affect classification performance.

The acquisition equipment included a finger clip with built-in sensors providing 3 physiological signals, i.e., heart rate (HR), skin resistance (SR), blood volume pressure (BVP), and 2 estimated parameters, namely the gradient of the skin resistance (GSR) and the speed of the changes in the data (CS - a measure of the signal’s’ entropy). Data is acquired at 15-16 samples per second (depending on the processor load) and filtered using a built-in Bessel filter. One of the main advantages of this type of sensing device is that it can be implemented as a wireless wearable equipment suitable for experiments inside IIE.

4.2 AANN Training

Training of the AANN was performed utilising the MATLAB implementation of the Levenberg-Marquardt algorithm in combination with Bayesian regularization (Foresee and Hagan, 1997) for enhanced generalization. Training data comprised the entirety of the neutral data from the eight individuals before physical activity totalling 18788 records. In order to reduce weight values and increase algorithm performance, data provided to the AANN was normalised to fall in the range between 0 and 1 using the maximal and minimal value of each physiological signal.

4.3 Attribute selection

With the purpose of selecting the physiological signals that provide the best class separation, the trained AANN was provided with a dataset containing both neutral and non-neutral data from the eight individuals. The resulting residual calculation was then subjected to a clustering analysis.

The calculation of the Davies-Bouldin Index (DBI) evidenced that the Heart Rate (HR) was the best attribute to distinguish between neutral and non-neutral data with 14.35 followed by Change Speed (CS) with 21.42. In the same manner, the CS was shown to be optimal for classifying positive and negative emotions with a DBI of 14.76 followed by the HR with 21.36.

4.4 SPRT implementation

It has been mentioned that the formulation of the SPRT is based on the PDF of each variable involved in the classification process. In order to safely assume that the residual values used in the SPRT calculation stem from a normal distribution and to justify the utilization of (1), it is necessary to assess the normality of the data. This can be done using what is referred to as a normality test using a normal probability plot. In a normality graph, the closest data points are to the straight line, the better the normality.

Figure 1 shows that the assumption of normality in the HR and CS after the elimination of the noise introduced by the sensing device is reasonable. For the purpose of guaranteeing an accurate detection of emotion states without compromising system response, a standard significance value of 0.05 for both alpha (α) and beta (β) was chosen thus proving a 99.5% confidence in the SPRT results (this value for α and beta β is similar to that typically used in statistical analysis). The mean (μ_1, μ_2) and variance (σ^2) values associated with the normal distribution of the residual were estimated from the AANN results to the entire dataset and were different for the SPRT implementation used for detecting affective status and that for emotional valence.

[Figure 1 goes here]

5. Results and Discussion

In order to test the learning and classification performance of the system, data randomly selected from an experimental session involving one of the original 8 volunteers and not used during training (i.e. after physical exertion) was provided to the AANN and SPRT modules (See Table 1). It is worth mentioning that this volunteer was among those regarded as high emotionally intense according to the AIM results.

[Table 1 goes here]

Results in Table 1 confirm that the trained AANN was capable of appropriately learning and identifying the relationships between sensors associated with the emotional episodes embedded in training physiological data.

Next, with the intention of determining whether the detection system was capable of generalizing from the original dataset and recognizing emotional states on any subject, an *entirely new* set of physiological data was collected from a physically fit and high-affect intensity female aged 48. Emotional states were elicited using 21 pictures from the International Affective Picture Systems (IAPS) (Center for the Study of Emotion and Attention, 1999) presented on the screen for 6000 ms with 35000 ms inter-slide blanks. The selection of the IAPS pictures was based on their high arousal and valence values in order to guarantee an optimal response from the participant.

Verbal self-reports after the presentation evidenced that some of the pictures did not evoke the targeted emotions in this particular subject. In order to identify data relating to the failed stimuli, a series of Wilcoxon similarity tests were performed on the 7 datasets associated with the positive and negative emotional episodes. The identification of failed stimuli using both subject accounts and mathematical tools instead of an entirely subjective discrimination enhances the validity to the results and provides more veracity to the analysis.

Table 2 shows the results obtained after the collected data were provided to the trained AANN.

[Table 2 goes here]

It can be seen from Table 2 that the emotional system was able to recognize the majority of the emotional episodes with 71% accuracy for the entire set of 21 emotional episodes and 80% if data from the failed stimuli is not considered. Emotional changes (neutral and non-neutral) were recognized in ten out of 14 occasions while a positive emotion was erroneously labelled as neutral 4 times and a neutral as negative in 2 instances. The recognition rates shown in Tables 1 and 2 demonstrate that the performance of the method described here is comparable to the best results achieved through off-line analysis.

It is important to mention that the fact that some emotional episodes were not recognized by this system or were attributed to a conflicting category, could be related to the personal interpretation of the pictures' emotional content rather than to the performance of the classification system itself. For instance, the fact that an emotional episode was recognized by the system as a negative one when the actual stimulus was neutral, does not entail an error in the system but instead the detection of the subject's actual reaction to that particular stimulus. Because the parameters used in the two SPRT modules (means and variances) to distinguish the different emotional states were based on the residual values from the entire group of 8 original volunteers they would remain valid for a population with similar emotional characteristics. Furthermore, the use of small alpha (α) and beta (β) values in the SPRT module, guarantees the accuracy of the verdict provided by the system regardless of the uncertainty introduced by personal interpretation of the stimuli. In other words, the fact that in both experiments participants were of high emotional intensity makes it more likely that the intended emotional states actually occurred, and because the system was trained to detect and classify the physiological response associated with such emotional states, the validity of the results presented above should be safe.

It is worth noting that one prerequisite for the optimal operation of the AANN is the 0-to-1 rescaling of real-time data using the individual maximal and minimal recorded values for each physiological signal during experimentation. In fact, the use of maximal and minimal values other than the ones related to the subject being examined could lead to very poor AANN estimations and consequently low recognition rates.

Although an initial collection of data with the sole purpose of determining maximal and minimal physiological values could seem as a limitation of the system, it actually is a normal tuning or pre-training process used in many agent configurations,

6. Conclusions

The present work demonstrates that real-time detection of emotional states in pervasive computing is feasible. It has been shown that the collection of emotional data from eight subjects with different affect intensity levels was sufficient to provide an AANN with enough information to generalise in the presence of data not included during training. The enhanced generalization capabilities of this system were confirmed by recognizing emotional states on a subject in real-time fashion with a 71.4 accuracy based on the original categorization of the emotional stimuli and 80% after the elimination of not evocative or contradictory stimuli.

Furthermore, the research in this paper confirms previous studies demonstrating that the use of a pre-trained AANN in conjunction with a SPRT-based decision module provides a mechanism that can reliable distinguish emotional states with high recognition rates. For the purpose of this analysis it has been shown that emotion detection could be done in principle for any subject as long as his/her maximal and minimal physiological response is known. The combination of the recognition system introduced in this paper along with portable sensing equipment would provide the mechanisms needed to integrate emotional states into decision systems in ubiquitous computing.

Future work involves an in-depth analysis of the significance of emotional information in the modelling of user needs in pervasive computing and IIE in particular. An IIE agent with emotion recognition capabilities could be employed to investigate how ambience conditions including device settings and weather conditions influence emotions and how the user responds or adjusts to changes in the environment.

Current experiments on affective agents are in progress inside the iDorm2. The iDorm2 has been developed at the University of Essex to investigate and compare various paradigms of intelligent agents designed to operate inside IIE (Callaghan 2001) (Callaghan 05). A wearable embodiment of the detection mechanism presented in this paper has also been built to allow free movement and comfort. It is argued that a potentially significant insight into the interrelation between emotions, environment and user actions could be achieved using affective software agents. It is worth noting that the embodiment of the emotion detection system in an affective agent preserves the advantages associated with emergent functionality of robustness, adaptability, scalability, and immediate, accurate response to changes in the environment

Acknowledgments

We are pleased to acknowledge the EU IST Disappearing Computer funding program which has provided support for the iDorm2 (intelligent dormitory), an experimental testbed for embedded-agents research. Enrique Leon also appreciates the support of the Mexican National Council for Science and Technology (CONACYT). We would also like to thank Robin Dowling for his enormous support in the implementation of the sensoring interface.

References

Anderson, K., McOwan, P., 2003. Real-time Emotion Recognition using Biologically Inspired Models, in: Proc. 4th International Conference on Audio and Video Based Biometric Person Authentication, Surrey, UK, pp. 119-127.

Ark, W., Dryer, D.C., Lu, D.J., 1999. The Emotion Mouse, in: Proc. of Human Computer Interaction 1999, Edinburgh, Scotland.

Avent, R.R., Chong, T., Neal, J.A., 1994. Machine Vision Recognition of Facial Affect Using Backpropagation Neural Networks, in: Proc. 16th Annual International Conference of the IEEE, Engineering Advances: New Opportunities for Biomedical Engineers, Baltimore, MD, pp.1364-1365.

Bartlett, M.S., Littlewort, G., Fasel, I., Movellan, J.R., 2003. Real-time face detection and expression recognition: Development and Application to Human-Computer Interaction, in: Proc. CVPR Workshop on Computer Vision and Pattern Recognition for Human-Computer Interaction, Madison, WI.
Brooks, R. A. "A Robust Layered Control System for a Mobile Robot", IEEE Journal of Robotics and Automation, Vol. 2, No. 1, March 1986, pp. 14–23; also MIT AI Memo 864, September 1985.
Callaghan V, Clark G, Colley M, Hagras H Chin JSY, Doctor F “Intelligent Inhabited Environments”, BT Technology Journal , Vol.22, No.3. Klywer Academic Publishers, Dordrecht, Netherlands, July 2004
Callaghan,V, Clarke,G.S., Colley,M.J., Hagras,H.A.K., 'A Soft-Computing DAI Architecture for Intelligent Buildings', Book Series: Studies in Fuzziness and Soft Computing, Vol. 75, Soft Computing Agents: New Trends for Designing Autonomous Systems. pp.221, 2001
Callaghan V, Colley M, Hagras H, Chin J, Doctor F, Clarke G, Programming iSpaces: A Tale of Two Paradigms, in book Intelligent Spaces: The Application of Pervasive ICT part of the series Computer Communications and Networks Steventon, Alan; Wright, Steve (Eds.) approx. 455 p. 162 illus. ISBN: 1-84628-002-8, Dec 2005
Davies, D.J., Bouldin, D., 1979. A Cluster Separation Measure. IEEE Trans. on Pattern Recognition and Machine Intelligence. PAMI-1(2).
De Silva, L.C., Ng, P.C., 1999. Bimodal Emotion Recognition, in: Proc. 1999 International Conference on Face and Gesture Recognition, Grenoble, France, pp. 332-335.
De Silva, L.C., Hui, S.C., 2003. Real-time facial feature extraction and emotion recognition, in: Proc. 2003 Joint Conference of the Fourth International Conference on Information, Communications and Signal Processing, 2003 and the Fourth Pacific Rim Conference on Multimedia, Singapore, Singapore pp. 1310-1314.
Fernandez, R., 1997. Stochastic Modelling of Physiological Signals with Hidden Markov Models: A Step Toward Frustration Detection in Human-Computer Interfaces. M.Sc. Thesis, Department of Computer Science, M. I. T., Cambridge, MA.
Foresee, F., Hagan, M., 1997. Gauss-Newton approximation to Bayesian learning, in: Proc. of the 1997 International Joint Conference on Neural Networks, Houston, TX, p.p. 1930-1935.
Fu, K.S., 1968. Sequential Methods in Pattern Recognition and Machine Learning. Academic Press, New York, NY.
Goleman, D., 1995. Emotional Intelligence. Why it can matter more than IQ. Bantam Books, New York.
Healey, J., 2000. Wearable and Automotive Systems for Affect Recognition from Physiology. M.Sc. Thesis, Department of Computer Science, M. I. T., Cambridge, MA.
Healey, J., Picard, R., Dabek, F., 1998. A new Affect-Perceiving Interface and its Application to Personalized Music Selection, in: Proc. of the 1998 Workshop on Perceptual User Interfaces, San Francisco, CA.
Heyward, V.H., 1998. Advanced Fitness Assessment and Exercise Prescription, Second ed., Human Kinetics, Champaign, IL.
Hines, W., Darryl, J., Uhrig, R., 1997. Use of Auto-Associative Neural Networks for Signal Validation, in: Proc. of Neural Network Applications ‘97, Marseille, France.
Inanoglu, Z., Caneel, R., 2005. Emotive Alert: HMM-Based Emotion Detection In Voicemail Messages, in: Proc. Intelligent User Interfaces (IUI 05), San Diego, CA.

Kim, K.H., Bang, S.W., Kim, S.R., 2002. Development of person-independent emotion recognition systems based on multiple physiological signals, in: Proc. of the Second Joint EMBS/BMES Conference, Houston, TX, pp. 50-51.
Kramer, M. A., 1992. Autoassociative Neural Networks. Computers and Chemical Engineering. 16(4), 313-328.
Lang, P.J., Bradley, M.M., Cuthbert, B.N., 2001. International Affective Picture System (IAPS): Intruction Manuela and affective ratings. Technical Report A-5, The Center for Research in Psychophysiology, University of Florida.
Larsen, R.J., Diener, E., Emmons, R.A., 1986. Affect Intensity and Reactions to Daily Life Events. Personality and Social Psychology, 51(4), 803-814.
Leon, E., Clarke, G., Callaghan, V., Sepulveda, F., 2004a. Real-time detection of Emotional Changes for Inhabited Environments. Journal of Computers & Graphics, Special Issue on Pervasive Computing and Ambient Intelligence – Mobility, Ubiquity and Wearables. 5(23), 635-642.
Leon, E., Clarke, G., Sepulveda, F., Callaghan, V., 2004b. Optimised attribute selection for emotion classification using physiological signals, in: Proc. of the 26th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, San Francisco, CA, pp. 184-187.
Leon, E., Clarke, G., Callaghan, V., 2005a. Towards A Robust Real-Time Emotion Detection System For Intelligent Buildings, in: Proc. IEE International Workshop on Intelligent Environments, Essex, UK, pp. 162-167.
Leon, E., Clarke, G., Sepulveda, F., Callaghan, V., 2005b (In Print). Real-time Physiological Emotion Detection Mechanisms: Effects of Exercise and Affect Intensity, in: Proc. of the 27th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, Shanghai, China.
Lisetti, C., Nasoz, F., LeRouge, C., Ozyer, O., Alvarez, K., 2003. Developing Multimodal Intelligent Affective Interfaces for Tele-Home Health Care. International Journal of Human-Computer Studies, Special Issue on Applications of Affective Computing in Human-Computer interaction. 59 (1-2), 245-255.
Lu, P., Zhang, M., Hsu, T., Zhang, J., 2001. An Evaluation of Engine Faults Diagnostics Using Artificial Neural Networks. Journal of Engineering for Gas Turbines and Power. 123(2), 340-346
Minsky M "The Society of Mind", Simon & Schuster Inc 1985 ISBN 0-671-63713-5.
Nasoz, F., Lisetti, C.L., Alvarez, K., Finkelstein, N., 2003. Emotion recognition from Physiological Signals for User Modelling of Affect, in: Proc. of the 3rd Workshop on Affective and Attitude user Modelling, Pittsburgh, PA.
Nicholson, J., Takahashi, K., Nakatsu, R., 2000. Emotion Recognition in Speech Using Neural Networks. Neural Computing & Applications. 9(4), 290-296.
Osgood, C.E., Suci, G.J., Tannenbaum, P.H. (1957). The measurement of meaning. University of Illinois Press, Chicago, IL.
Picard, R., Vyzaz, E., Healey, J., 2001. Toward Machine Emotional Intelligence: Analysis of Affective Physiological State. IEEE Trans. on Pattern Analysis and Machine Intelligence. 23 (10), 1175-1191.
Rosenblum, M., Yacoob, Y., Davis, L., 1994. Human Emotion Recognition from Motion Using a Radial Basis Function Network Architecture, in: Proc. 1994 Workshop on Motion of Nonrigid and Articulated Objects, Austin, TX, pp. 43-49.
Salovey, P., Mayer, J. D., 1990. Emotional intelligence. Imagination, Cognition, and Personality. 9, 185-211.

Sepulveda, F., Meckes, M., Conway, B.A., 2004. Cluster Separation Index Suggests Usefulness of Non-Motor EEG Channels in Detecting Wrist Movement Direction Intention, in: Proc. of the 2004 IEEE Conference on Cybernetic and Intelligent Systems.

Smith, J.C., O'Connor, P.J., Crabbe, J.B., Dishman, R.K., 2002. Emotional Responsiveness after low-and moderate-intensity Exercise and Seated Rest. Medicine and Science in Sports Exercise. 34(7), 1158-1167.

Steels, L., 1991. Towards a Theory of Emergent Functionality, in: Proc. of the First International Conference on Simulation of Adaptive behaviour, Paris, France, pp. 451-461.

Sun, Y., Sebe, N., Lew, M.S., Gevers, T., 2004. Authentic Emotion Detection in Real-Time Video, in: Proc. Computer Vision in Human-Computer Interaction, ECCV 2004 Workshop on HCI, Prague, Czech Republic.

Tsuyoshi, M., Shinji, O., 1999. Emotion Recognition and Synthesis System on Speech, in: Proc. 1999 IEEE International Conference on Multimedia Computing and Systems, Firenze, Italy, pp. 840-844.

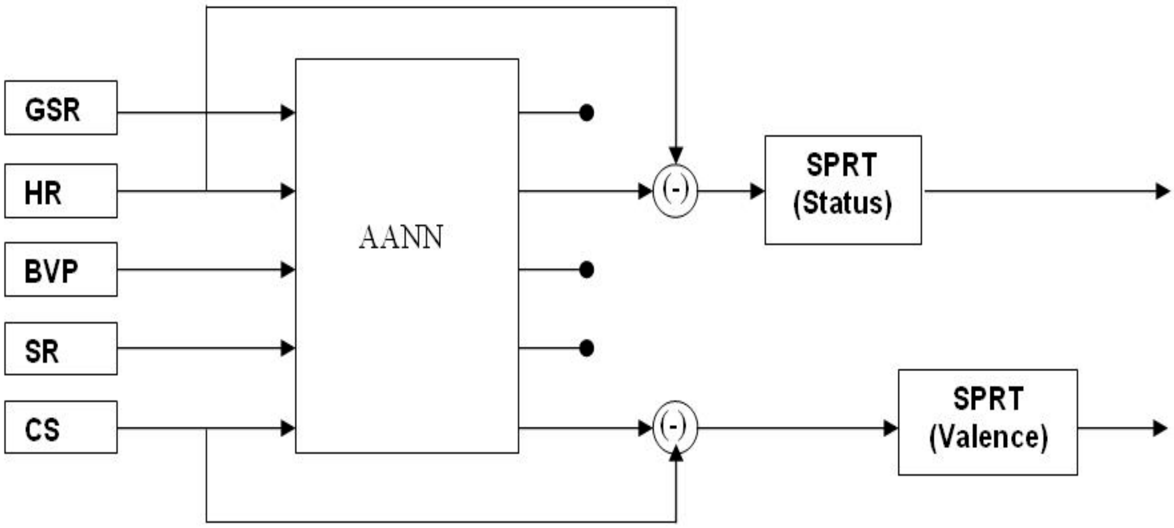

Figure 1. Diagram of the emotion recognition system. “(-)” denotes subtraction.

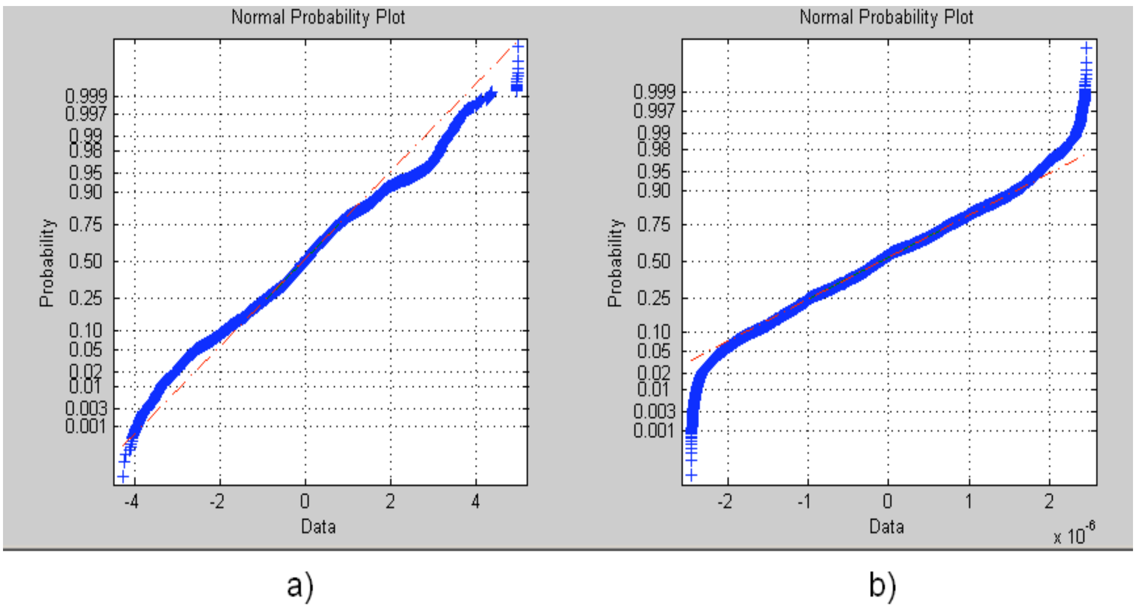

Figure 2. Normality test for the residuals produced by the AANN for a) HR and b) CS.

Emotional Stimuli (Original set)	Detection Output
Neutral	Neutral
Positive	Neutral
Positive	Neutral
Negative	Negative
Negative	Negative
Neutral	Neutral
Neutral	Neutral
Positive	Neutral
Neutral	Neutral
Negative	Negative
Positive	Positive
Negative	Negative
Positive	Positive
Neutral	Neutral
Negative	Negative
Negative	Negative
Positive	Positive
Positive	Positive
Neutral	Neutral
Negative	Negative
Neutral	Neutral
Overall recognition rate: 85.7%	

Table 1. Recognition results for 21 emotional episodes.

Emotional Stimuli (Original set)	Detection Output	Emotional Stimuli (Revised set)	Detection Output
Neutral	Neutral	Neutral	Neutral
Positive	Neutral	Neutral	Neutral
Positive	Neutral	Neutral	Neutral
Neutral	Neutral	Neutral	Negative
Neutral	Neutral	Negative	Negative
Negative	Negative	Positive	Positive
Neutral	Negative	Negative	Negative
Negative	Negative	Negative	Negative
Positive	Positive	Neutral	Neutral
Negative	Negative	Negative	Negative
Negative	Negative	Positive	Neutral
Neutral	Neutral	Negative	Negative
Negative	Negative	Neutral	Negative
Positive	Neutral	Neutral	Neutral
Positive	Neutral	Positive	Positive
Negative	Negative		
Neutral	Negative		
Neutral	Neutral		
Positive	Positive		
Negative	Negative		
Positive	Positive		
Overall recognition rate: 71.42%		Overall recognition rate: 80%	

Table 2. Recognition results for the original 21 emotional episodes and after the elimination of data corresponding to failed stimuli.

