User Model User-Adap Inter (2008) 18:1–9
DOI 10.1007/s11257-007-9044-7

PREFACE

# Introduction to special Issue on 'Affective modeling and adaptation'

**Sandra Carberry · Fiorella de Rosis**

Received: 14 October 2007 / Accepted: 14 October 2007 / Published online: 17 January 2008
© Springer Science+Business Media B.V. 2008

> *We are all ruled in what we do by impulses; and these impulses are so organized that our actions in general serve for our self preservation and that of the race. Hunger, love, pain, fear are some of those inner forces which rule the individual's instinct for self preservation. At the same time, as social beings, we are moved in the relations with our fellow beings by such feelings as sympathy, pride, hate, need for power, pity and so on.*
>
> Albert Einstein, 1950

## 1 The framework

Emotion is generally regarded as an involuntary mental response, such as surprise or anger, that is accompanied by physiological changes. Ekman was one of the early pioneers in characterizing facial expressions that convey human emotion, and Ekman and Friesen argued that there are six basic emotions that are exhibited in human facial expressions: anger, disgust, fear, happiness, sadness, and surprise (Ekman and Friesen 1978). However, in recent years, researchers have expanded the range of psychological conditions of interest, and have used the term *affect* to refer to elements of this larger set of states, not all of which would qualify as actual emotions.

S. Carberry
Department of Computer Science, University of Delaware, Newark, DE 19716, USA
e-mail: carberry@cis.udel.edu

F. de Rosis (✉)
Department of Informatics, University of Bari, Via Orabona 4, 70126 Bari, Italy
e-mail: derosis@di.uniba.it

The study of affective states has been motivated by specific applications which presumably would be more successful if the system could identify that the user was exhibiting such a state and could adapt the system's response accordingly. For example, overall learning rate might be improved if a tutoring system could detect dwindling student interest in the tutoring lesson and adapt its interaction so as to increase the student's interest level. Thus the range of affective states that have been studied includes both traditional emotions such as happiness and surprise as well as other states such as boredom, confusion, frustration, and uncertainty.

Affective computing may be subdivided into four areas. The first is the analysis and characterization of affective states, empirical experiments that identify states that are exhibited in natural interactions, and an analysis of the relationship between affect and cognitive processes such as learning. The second is the automatic recognition of affective state. In addition to analyzing facial expressions, research has extracted cues to affective state from linguistic expression, acoustic signals such as prosody, posture, head nods, eye gaze, and physiological responses such as heart rate. The third area is adapting the system's response to the user's particular affective state. And lastly, affective computing includes the design of avatars that themselves exhibit appropriate affective state, with the goal being lifelike characters and more effective interactions.

This special issue focuses on the first two areas of affective computing. Although there has been much research on lifelike avatars and some work on adaptation based on the user's affective state, the success of these areas is heavily dependent on being able to accurately recognize the affective state of the user. Thus we regard the analysis and modeling of user affect as fundamental and critical areas of research.

The modeling of affect has been pursued in very different domains. Tutorial domains have been the most popular due to the great interest during the past two decades on individualizing learning via automated tutoring systems and the existence of working tutoring systems for many applications. However, modeling affect is also important in diverse areas such as medical consultation, help systems, interactive games, and even potential driver monitoring systems in automobiles.

The reasoning methods used in modeling affect have also varied. Some early work utilized rules based on linguistic expression, context, and stereotypical beliefs to recognize affect from a user's utterance. However, in recent years, the emphasis has been on machine learning and statistical approaches that learn models for identifying affect from a corpus of examples. Furthermore, researchers have placed greater emphasis on critically evaluating their models.

## 2 Papers in this issue

The papers in this special issue represent work from some of the top research groups on affective modeling. As would be expected given the preponderance of work on tutoring systems, four of the papers focus on characterizing and modeling affect during a tutorial interaction. However, they differ in a number of ways.

The work by Forbes-Riley, Rotaru, and Litman is part of the ITSPOKE project in which students interact via spoken language dialogue with a physics tutor. Their paper explores the relationship between affect and learning. They consider only two kinds of affect: frustration and uncertainty. These were selected because they are prevalent in tutoring interactions. In addition to the affect parameters, they also consider what they term system-generic features such as average student words per turn, speech recognition features such as word error rate, features related to the correctness of the student's answers, and features reflecting discourse structure transitions. They then use multivariate regression to develop models of student learning, both for single features and for bigrams. Their results show that the affective states of frustration and uncertainty correlate strongly with student learning, as do several discourse structure transition features and bigrams representing combinations of these individual features. They hypothesize that the correlations of affect with learning results from students exhibiting frustration and/or uncertainty when they are actively engaged in the learning process. Their further experiments showed that adding affective features to learning models generally improves the quality of the model. Since the work by Forbes-Riley et al. shows that student affect correlates with learning, it suggests that tutoring systems can use student affect to hypothesize whether learning is occurring and to adapt their strategies accordingly.

The paper by D'Mello, Craig, Witherspoon, McDaniel, and Graesser is part of the Auto Tutor project in which student responses consist of typed natural language input. Their paper investigates the relationship between affect and dialogue features (such as the number of words in a student's response, elapsed time of the tutoring session, time between tutor question and student response, and the type of speech act performed by the student), and whether affect can be recognized from such dialogue features. Whereas Forbes-Riley et al. studied only frustration and uncertainty, D'Mello et al. consider frustration and five additional affective states: boredom, confusion, delight, flow, and surprise. The paper considers many different modeling techniques, including regression analysis, Bayesian classifiers, nearest neighbor, decision trees, and support vector machines. Their results show a correlation between dialogue features and affective state and moderate success in recognizing affect via automatically learned classifiers.

McQuiggan, Mott, and Lester also investigate the modeling of affect. They consider the recognition of the student's level of self-efficacy or confidence in their ability to succeed in a given situation. Their work takes place in the context of CRYSTAL ISLAND, an interactive inquiry-based learning environment for exploratory learning of genetics. The paper investigates predicting self-efficacy from four kinds of features in addition to demographic data: temporal features such as the elapsed time since the student achieved a goal, locational features that capture facets such as the student's location on the imaginary island (and thus whether the student is in a situation where learning tasks can be performed) and the number of times a student has visited a particular location, intentional features such as the rate of achieving goals, and physiological readings of galvanic skin response and blood volume pulse (from which heart rate was computed) provided by a biofeedback device on the student's hand. Initial experiments in an online tutorial system showed that physiological response data was important in predicting self-efficacy. Thus experiments in the inquiry-based learning environment incorporated the physiological characteristics. Using both Naïve Bayes and decision tree methods, McQuiggan et al. constructed models of self-efficacy that were significantly better in the complex inquiry-based learning environment than baseline models.

The aim of the research by Porayska-Pomsta, Mavrikis and Pain is to introduce adaptation to some affective factors (confidence, interest, effort) by reproducing the criteria applied by tutors. Toward this goal, they collect a corpus of computer-mediated student–tutor interactions and ask tutors to annotate these dialogues before they are processed statistically (via decision trees). What is notable, in particular, about this paper is that not only is a recognition method proposed, but also a set of adaptation rules are produced, which are the first step towards defining how adaptation to affective factors might be realized.

Although tutoring has been the most popular application domain for affective computing research, many other domains presumably would benefit from modeling affect. The next two papers consider different application domains than ITSs, both oriented to children as potential users. In both cases, research is focused on recognition of affective factors, and does not yet appear to be at a stage of proposing criteria for adaptation.

Batliner, Steidl, Hacker and Noth's research is aimed at adapting the behavior of a robot to the attitude displayed by users who communicate with the robot via spoken commands. The data they consider consist of children's directive commands to a commercial robot (AIBO). Their data shows that commands may be correlated with different kinds of emotions: being angry or joyful, versus motherese or reprimanding. The corpus of human–robot interactions considered in the paper was collected with a Wizard of Oz study, a method that is applied quite frequently in studying communication with artificial agents. Corpus annotation was performed by five raters using majority agreement to define an external reference; recognition is based on a combination of linguistic and acoustic features, and nonmetrical multidimensional scaling is applied to process the data. The resulting solution models valence as the first dimension and social interaction as the second.

Yannakakis, Hallam and Hauptop Lund consider interaction with a tangible game (Playground) as a tool to study to what extent fun can be recognized from integration of physiological signals (heart rate) with self-assessment data. The final aim of this research is to estimate the degree to which games engage children, so as to adjust digital entertainment environments to their preferences. Data analysis methods applied are, in this case, a combination of regression models, neural networks and genetic algorithms.

## 3 Perspectives

In addition to intelligent tutoring and game playing (applications that are represented in this Special Issue), affective computing offers promise for improving performance in other areas: examples include support for car drivers who exhibit stress or fatigue, smart houses with the ability to monitor and provide assistance in emotional situations, assistance for special categories of users such as the elderly or autistic children at school (El Kaliouby et al. 2006), and interactive digital TV. The Proceedings of the 2007 Conference on Affective Computing and Intelligent Interfaces are a good reference point for an updated view of the present situation (http://www.gaips.inesc-id.pt/acii2007/).

The interaction modality that is particularly expected to profit from the ability to recognize, process, and appropriately exhibit affect is that of artificial agents, be they embodied characters or robots (see Trappl et al. 2002, and innumerable examples in the Proceedings of the Conferences IVA 2007 and RO-MAN 2007). However, the applications mentioned above go beyond this modality, to consider forms of interaction of more general use, such as mobile systems, natural language dialogues, or special-purpose graphical interfaces.

Are we close to the goal of building a system that is able to adapt to affective factors? Probably not: to our knowledge, no such system is yet in use, and probably none of the existing models has yet advanced to the stage where actual deployment is on the horizon. Several methodological problems must be solved, before such a system becomes a reality. The following summarizes what we consider to be the main obstacles that must be overcome.

### 3.1 Recognizing the affective state of the user

This is the domain in which research work is focused at present. Information sources considered for monitoring changes in affective states range from biophysical signal processing (such as in Mc Quiggan et al and Yannakakis et al. in this Issue) to speech analysis (such as in Batliner et al. this Issue), and observation of face, gesture, and body posture. Recently, research on emotion recognition from speech has begun to place greater emphasis on the actual language used rather than being concerned primarily with the features of the acoustic signal, as evidenced in the papers by Forbes et al. Porayska-Pomsta et al. and D'Mello et al.

Good recognition rates have been obtained for various classes of emotions, by using not only artificial acted data (as was done in the early days of affective computing research) but also natural data which, though more difficult to collect and recognize, promises models that better capture reality. Relatively stable affective factors (such as personality traits) have been recognized with classical questionnaires, such as Myers-Briggs for the Big-Five classification (McCrae and Costa 1987), or with, again, language analysis methods (Gill and Oberlander 2002); these represent respectively explicit acquisition and implicit acquisition. A parallel research vein at MIT[1] is exploring the idea of building new interaction devices, such as an emotional mouse or a touch-sensitive seat, to correlate user actions with their emotional state.

### 3.2 Integrating affective states into consistent user models

Affective states of users' are known to influence the rest of their minds. According to various psychologists, emotions influence beliefs and goals and are in turn influenced by them (Oatley and Johnson-Laird 1987; Frijda et al. 2000). Interests, preferences

1 http://www.affect.media.mit.edu/projects.php?id=1104

and motivations are influenced as well, either permanently or temporarily, by both stable and unstable affective factors. A few authors even tend to consider some of these features in the category of affective factors, which is probably not correct, or at least does not fit with the HUMAINE project's proposal for systematizing concepts and terms in affective computing.[2]

Simply recognizing an emotion is not sufficient to enact an effective adaptation process. After recognizing with a high level of accuracy that a user is fearful, angry, bothered, or frustrated, the application must try to infer the reasons for this emotional state. It must consider the context in which the emotion was recognized and integrate the affective and non-affective components of the user model that will drive adaptation. The application therefore needs to build and update dynamically an integrated 'rational&emotional' model, by dealing with the innumerable sources of uncertainty that derive both from the limited accuracy of the recognition process and from the way affective and rational factors influence each other. Cognitive emotion models with different knowledge grain sizes have been proposed; these include EMA (Marsella and Gratch 2006), the model in (Conati and McLaren 2005), and Emotional-Mind (Carofiglio et al. in press). These and other models tend to rely on the famous Ortony, Clore and Collin's psychological theory (Ortony et al. 1988). The papers in this special issue consider other emotions (such as frustration, enthusiasm, anxiety, etc) that are common and relevant in human–computer interaction.

Affective user modeling is a very rich and fertile domain, especially with respect to integration of emotional and rational aspects of behavior. Humans do not seem to be consistent in their emotional and rational thinking and behavior; on the contrary, cognitive dissonance (Festinger 1957; Harmon-Jones and Mills 1999) and misleading emotions (Goldie 2000; Gigerenzer and Goldstein 1996) seem to occur quite frequently in real life. This raises the prospect of reconsidering the concept of consistency in user models and dealing with recalcitrant emotions appropriately.

### 3.3 Responding to user affect

Several kinds of actions can follow recognition and modeling of the user's affective state: emotions may be regulated, when their impact might be harmful or in any way detrimental to the user (such as stress during vehicle driving), or strengthened when they are expected to have a positive impact on the user's performance (such as enthusiasm in tutoring systems). In other situations, knowledge of the user's affective state can be employed to display empathy or, more generally, to harmonize an artificial agent's behavior with the user's behavior; a typical example is the triggering of 'small talk' when interacting with socially warm users (Bickmore and Cassell 2005; de Rosis et al. 2006).

More ambitiously, the system might attempt to elicit certain emotions in order to enhance effectiveness of the application. In this case, the idea is to exploit emotions as support for achieving the application goal: such emotions include enthusiasm in tutoring sessions, fun or enjoyment in game-playing, and even emotions such as fear

[2] http://www.emotion-research.net/ws/conceptualizingemotion/

or the anticipation of pleasure in persuasive dialogues (Walton 1992; O'Keefe 2002; Miceli et al. 2006). However, the emotions that will be induced by a given communicative action can only be forecast with a margin of uncertainty, as they are highly dependent on the context, the user personality, his or her background etc; a plausible model of the user is thus essential in order to respond to user affect appropriately and to repair possible errors.

This is the area in which affective computing research is still in its infancy. To our knowledge, systems capable of adapting to the user's affective state have not yet been developed. One reason for this void might be the relative paucity of psychological studies that clearly specify the impact of different emotional states (ones that typically occur during human–machine interaction) on reasoning and behavior in humans, and how this influence may be enhanced or mitigated. In other words, knowledge about the cycle of emotion recognition, its influence on action, and the resulting mental changes is still incomplete; further elucidation and clarification requires interdisciplinary collaboration with psychologists and philosophers.

### 3.4 Evaluating affective systems

This is a particularly problematic aspect of research on affective computing. Assessing the accuracy of models of emotion recognition requires an objective reference. Subjective assessment with questionnaires based on Likert scales is far from ideal, due to the difficulty that humans have in identifying their own emotional state. More immediate and natural assessment methods have been proposed, such as haptic interaction tools (Picard and Daily 2005; Höök et al. 2007). For this reason, definition of a markup language and the labeling of data by a group of external raters, so that the data can be used to train the recognition model, is usually the first phase of any recognition study. (This is the approach that is also taken by the papers in this Special Issue.) In evaluating affective systems, such as an emotional embodied conversational agent, a black box approach is employed, for instance by using questionnaires to compare affective with non affective versions of the same application with respect to criteria of plausibility, acceptability, likeability, etc. This method is not fully satisfactory, as it does not facilitate identification of the reasons for possible malfunctions (that is, whether the cause of the malfunction is in the affect recognition and modeling component or in the adaptation criteria).

*To conclude*: we claim much research is still necessary before we will see systems that adapt effectively to the affective state of the user, be it an emotional state, an attitude, or a personality trait. However, interest in this intriguing domain has increased considerably over the last 10 years. Evidence of this expanding interest is the last book by Marvin Minsky, which proposes the following research agenda:

> " ... many thinkers still insist that machines can never feel or think... That once was a popular belief, but today it is widely recognized that behavior of a complex machine depends only on how its parts interact, but not on the 'stuff' of which they are made... This suggests replacing old questions like "What sorts of things are emotions and thoughts?" by more constructive ones like "What

processes does each emotion involve?" and "How could machines perform such processes? (Minsky 2006)

The birth of an International Association on Emotions (HUMAINE) is, at the same time, proof of this interest and a forum for encouraging future interdisciplinary cooperation in the area. If this interest continues to grow, functioning affective systems might become a reality earlier than we currently envision.

## References

Bickmore, T., Cassell, J.: Social dialogue with embodied conversational agents. In: van Kuppevelt, J., Dybkjaer, L., Bernsen, N. (eds.) Advances in Natural, Multimodal Dialogue Systems, pp. 1–32. Kluwer Academic, New York (2005)

Carofiglio, V., de Rosis, F., Grassano, R.: Dynamic models of mixed emotion activation. In: Canamero, L., Aylett, R. (eds.) Animating Expressive Characters for Social Interactions. John Benjamins Publ Co., Amsterdam (in press)

Conati, C., McLaren, H.: Data-driven refinement of a probabilistic model of user affect. In: Tenth International Conference on User Modeling, pp. 40–49. Edinburgh, Scotland (2005)

de Rosis, F., Novielli, N., Carofiglio, V., Cavalluzzi, A., De Carolis, B.: User modeling and adaptation in health promotion dialogs with an animated character. J. Biomed. Inform. **39**(5), 514–531 (2006)

Ekman, P., Friesen, W.V.: The Facial Action Coding System: A Technique for the Measurement of Facial Movement. Consulting Psychologists Press, Palo Alto (1978)

El Kaliouby, R., Picard, R.W., Baron-Cohen, S.: Affective computing and autism. Ann. N.Y. Acad. Sci. **1093**(1), 228–248 (2006)

Festinger, L.: A Theory of Cognitive Dissonance. Stanford University Press, Stanford (1957)

Frijda, N.H., Manstead, A.S.R., Bem, S.: Emotions and Beliefs: How Feelings Influence Thoughts. Cambridge University Press, Cambridge (2000)

Gigerenzer, G., Goldstein, D.G.: Reasoning the fast and frugal way. Psychol. Rev. **103**, 650–669 (1996)

Gill, A.J., Oberlander, J.: Taking care of the linguistic features of extraversion. 24th Annual Conference of the Cognitive Science Society, pp. 363–368. Washington, DC, (2002)

Goldie, P.: The Emotions. Oxford University Press, Oxford (2000)

Harmon-Jones, E., Mills, J.: Cognitive Dissonance: Progress on a Pivotal Theory in Social Psychology. American Psychological Association, Washington (1999)

Höök, K., Isbister, K., Laaksolahti, J., Sharp, M.: Sensual evaluation instrument: developing a trans-cultural self-report measure of affect. Int. J. Hum. -Comput. Stud. **65**(4), 315–328 (2007)

Marsella, S., Gratch, J.: EMA: a computational model of appraisal dynamics. In: Gratch, J., Marsella, S., Petta, P. (eds.) Agent Construction and Emotions, pp. 601–606. Austrian Society for Cybernetic Studies, Vienna (2006)

McCrae, R., Costa, P.T.: Validation of the five-factor model of personality across instruments and observers. J. Pers. Soc. Psychol. **52**(1), 81–90 (1987)

Miceli, M., de Rosis, F., Poggi, I.: Emotional and non-emotional persuasion. Appl. Artif. Intell. **20**, 849–879 (2006)

Minsky, M.: The Emotion Macchine. Simon and Schuster, London (2006)

Oatley, K., Johnson-Laird, P.N.: Towards a cognitive theory of emotions. Cogn. Emot. **1**(1), 29–50 (1987)

O'Keefe, D.J.: Persuasion, Theory & Research. Sage Publications, London (2002)

Ortony, A., Clore, G.L., Collins, A.: The Cognitive Structure of Emotions. Cambridge University Press, Cambridge (1988)

Picard, R.W., Daily, S.B. Evaluating affective interactions: alternatives to asking what users feel. CHI Workshop on Evaluating Affective Interfaces, Portland, Oregon. http://www.affect.media.mit.edu/pdfs/05.picard-daily.pdf

Trappl, R., Petta, P., Payr, S.: Emotions in Humans and Artifacts. The MIT Press, Cambridge (2002)

Walton, D.N.: The Place of Emotions in Argument. The Pennsylvania State University Press, University Park (1992)

**Authors' vitae**

**Sandra Carberry** is a Professor of Computer Science at the University of Delaware. She is a member of the editorial board of User Modeling and User-Adapted Interaction and served for 6 years as an officer of the Association for Computational Linguistics. Her research interests are in the areas of discourse and dialogue, with a focus on intelligent interfaces, affective computing, and user modeling.

**Fiorella de Rosis** is a Professor of Computer Science in the Department of Informatics, University of Bari. She is a Member of ACM and ISRE (the International Society for Research on Emotions) and a member of the Executive Committee of HUMAINE (the International Association on Emotions). Her research interests are in the area of intelligent interfaces, user modeling and adaptation, and affective computing.