Preface:

Towards Adaptation of Interaction to Affective Factors

FIORELLA DE ROSIS

The recent increase of interest in ‘affective computing’ is probably due to two—apparently opposite—trends in HCI. On the one hand, the shift of hardware towards computers that are, in a way, ‘closer’ to the Users body, by being either wearable or ubiquitous, opens the opportunity to monitor new parameters about the Users state and to send them signals that they may seize with all their senses. On the other hand, the shifting of HCI metaphor from ‘using a tool’ to ‘interacting with a friend’ that takes the appearance of an Embodied Animated Agent that ‘gives the illusion of life’ requires recognising the User’s affective state and reacting appropriately to this state, by showing an affectively compatible behavior (Bates, 1994).

Reflections on how affective factors might influence human–computer interaction are rather recent. The major experimental background to this relatively new research domain was given by the well-known experimental studies at the University of Stanford, which proved that ‘*people respond to computer personalities in the same way they would respond to human personalities*’ (Nass et al., 1995; Moon & Nass, 1996; Isbister & Nass, 1999). The more general claim that ‘humans respond socially to technology’ was supported by other studies from the same Department, which showed that people tend to attribute ‘responsibilities’ to computers, are susceptible to their flattery and tend to form team relationships with them (Nass et al., 1996; Fogg & Nass, 1997; Moon & Nass, 1998). In these studies, the subjects’ personality was evaluated in terms of the ‘Five Factor Model’ (FFM: McCrae & John, 1988), by administering questionnaires that are of common use in psychological research. The interface personality was also labelled in terms of the FFM: in particular, Extraversion (or Dominance) and Agreeableness. The results of studies by Nass and colleagues were acknowledged as a real novelty in the HCI domain, so that they became the main reference point to the computer scientists who intended to enrich interfaces with ‘some personality trait’: FFM is now the theory behind the majority of these systems: see, for instance, (André et al., 1999; Arafa et al., 1999; Taylor et al., 1999).

Other Projects were aimed at endowing the interface with ‘emotions’ rather than with a ‘personality’¹: this happened, in particular, in educational systems, in which showing ‘empathy’ with the emotional state of the student is considered a property that increases the system’s chance of success (just an example: Lester et al., 2000). In defining how to endow a system with an emotional state, the main reference is Ortony, Clore and Collins’ theory (OCC: Ortony et al, 1988) and its further developments by Elliott (Elliott & Siegle, 1993).

Finally, some authors studied how other attitude-related variables, like ‘cognitive overload’ or ‘doubt’, may be recognised through speech or language style analysis (Berthold & Jameson, 1999; Carberry & Schroeder, 2001) or how simulated dialogs might be enriched with affect-related issues such as ‘humor’, ‘deception’ or ‘politeness’ (Elliott, 1999; Castelfranchi et al., submitted; Ardissono et al., 1999).

However, the first attempts to model affective factors and to adapt human–computer interaction to some of them were performed about 20 years ago: to make just one remarkable example, the PhD Thesis of Carbonell was about how some knowledge of personality could improve story understanding. In a review of these systems, Pfeifer distinguishes between systems aimed at ‘augmenting AI models of cognition’ and ‘AI models of emotion’. In the first category, he places the models in which ‘cognition is, in a sense, augmented by emotion’ and in the second, those in which ‘emotion is the central focus’ (Pfeifer, 1988). The first experiences reviewed in this paper date back to the beginning of years sixties, although the majority of them concentrates from the years eighties on; the prevailing application domain is conversation and natural language processing.

Why, then, after so long is this field still seen as ‘something far away in the future’, a goal difficult to achieve? How far are we from realising an interface that really adapts to the ‘affective’ state of the User, even if only in specific and well delimited domains? This Special Issue is aimed at promoting a reflection on this subject in our Journal, with a set of papers that refer to significant experiences in the domain. It comes after two Workshops on ‘Attitudes, Personality and Emotions in User-Adapted Interaction’ which were organised in the scope of the ‘User Modeling’ 1999 and 2001 Conferences (contributions to these Workshops may be found at the websites: <http://aos2.uniba.it:8080/ws-um99.html> and <http://aos2.uniba.it:8080/ws-um01.html>). Let me, first of all, briefly present these papers.

¹Though not being an expert in cognitive psychology, I believe that there is sometimes an inappropriate use of these terms in papers dealing with ‘believable agents’, and a tendency to mix up personality traits with emotions. As far as I know, only some personality traits characterise a *propensity for a given emotional reaction*, while, in general, they denote different and much more stable characteristics than emotions: see, for instance, Bianchi-Bertouze and Lisetti, in UMUI 12 (1), 2002, but also Castelfranchi (2000).

1. The Papers in this Special Issue

Although the Call was open to all problems related to adapting User Interfaces to 'affective' factors, the contributions that appear in this Special Issue² are focused on emotions and the way they may be recognised, manifested and described: these methods are essential to build systems that are able to speculate about the emotional state of a User and to respond in a way that is appropriate to the situation, irrespective of whether they have a graphical, a natural language or an agent-based interface.

The first two papers (by Bartneck & Mozziconacci) deal with the problem of how some emotional states that are described in the literature (for instance, in the OCC theory) may be manifested through facial expressions or through the voice. Bartneck's contribution is aimed at responding to the question of whether a refined representation of the human face is needed to manifest these basic emotions: this question is relevant, considering the evolution of technology towards interfaces of small dimension and with a low graphical resolution, in computers with limited memory and software resources. The author compares, in particular, the 'convincingness' of a face that is rendered with decreasing levels of graphical complexity with that of a 'natural human' face or of a 3D 'computer-rendered' face a very simple bidimensional image accompanied by an audio. He describes a set of experimental studies which show that some emotions (fear, for instance) are more difficult to represent and less convincing than others (surprise and happiness, for instance). The *convincingness* of emotions does not seem to be influenced significantly by abstraction, while it seems to be favoured by the combination of different media (audio and video). Another result of Bartneck's studies that merits reflection is that the human ability to distinguish among the different facial expressions that were shown to them does not guarantee that the emotion they attribute to these faces are the same as those that the study designers intended to attribute them. On one side, these results are encouraging, as they seem to prove that a sophisticated technology is not really necessary to represent emotions in embodied animated agents. On another side, they are not reassuring because, as the author says, *'user-machine communication might fail if the expression is frequently misinterpreted'*.

The results of the studies about emotional speech synthesis that are described in the paper by Mozziconacci seem to head in the same direction. In these studies, natural speech was compared with its manipulation via analysis-resynthesis and with synthetic speech. Also in this case, some emotions seem to be more easily recognisable than others (again, fear, for instance); again, the cultural context and the situation can influence the human interpretation of signals. But the positive finding is that, by manipulating the values of only three acoustical parameters (pitch level, pitch range and speech rate), different emotions, such as boredom, sadness or indignation may apparently be represented in a rather 'believable' way.

²The last two papers plus the book reviews will appear in UMUI 12 (1), 2002.

A common question that is unanswered by the two studies is how much the emotion *intensity* affects its perception and the convincingness of its representation; that is, whether exaggerated emotional expressions (in both the face and the voice) may be avoided without reducing correct recognition. Another unsolved question also remains: according to Moon and Nass, to answer positively to the question: '*Can computer personalities be human personalities?*', computer systems should satisfy not only the 'identification requirement' (individuals should correctly *identify* the particular personality that is manifested in the computer) but also the 'behavioral requirement' (individuals should *behave*, towards the computer, in the same way they would behave towards a person with the same personality) (Moon and Nass, 1996). This test appears to be even more difficult to pass than the famous Turing's imitation game (Turing, 1950).

Should a machine, though, have a face and a voice to put empathy in its interaction with the Users and to accord its behavior to their affective state? Certainly not: innumerable proofs in favour of this claim have already been provided in the studies by Nass and colleagues, which proved that even common interfaces, in which only the name given to the computer, the phrasing of the text displayed, the interaction order and the graphical aspects were manipulated, could show personality traits and emotions that might be interpreted in terms of the FFM. Similar results were shown in studies concerning natural language communication, in which the discourse structure and the argumentation style were varied according to the emotional state of the participants in the dialog and their relationships (Wegman, 1988; Loyall & Bates, 1997; Walker et al., 1997; de Rosis et al., 1999, 2000).

An example of a graphical interface that adapts to the emotional state of the User is described in Hudlicka's paper. The system described (ABAIS) reacts to the User's anxiety state by varying the interface layout and the type of information displayed. The underlying hypothesis is that affective states have a major impact on performance (a problem that is especially serious in the application domain considered, that of air pilots) and that the consequences of some of these effects can be predicted: in particular, they alter the nature of attentional processing, activate or inhibit perceptual and cognitive schemata and promote or inhibit the selection of particular actions. The User's anxiety state is assessed via a knowledge-based approach which combines a variety of user data (e.g. personality traits, individual history, simulated physiological data) and task factors to derive the most likely user state; the system predicts the impact of this state on the user behavior and selects a strategy that aims at avoiding possible errors through an adaptation of the interface. The approach was implemented and demonstrated in a prototype of a fighter pilot simulation and demonstrated the feasibility of adapting to the affective state of the User, although only for a specific state (anxiety), via a specific method (GUI format), and for a specific category of Users (air pilots). The architecture and principles behind its design provide, however, useful hints on how systems with adaptation abilities might be built, in similar contexts.

Are we sure, though, that we know exactly what we are talking about? Are we certain that the terms which are commonly employed to describe emotional states or visual impressions correspond to clear, time invariant and sharable concepts? This is not a minor question because, if it is not so, the assumptions behind the design of any system that is able to recognise and manifest affective states are not satisfied. '*Of all human experiences, emotions seem to be among the least susceptible to verbal description*' (Clore and Ortony, 1988). This basic question is the common core of the two systems developed by Lisetti and Bianchi-Bertouze (MOUE and MIKE) that are described in the last paper of this Issue. At the basis of MOUE, there is a theory of emotion concepts as an ontology whose nodes are classical schemata of Artificial Intelligence; every emotion is described through a list of *components* (valence, intensity, duration, focality, etc). Rather than attempting a systematic description of the main affective states (which are usually denoted with labels like 'happy', 'sad', 'surprised', etc), MOUE intends to provide a tool to build these descriptions as 'computational schemata', through a system-User interaction. The User's face expression is captured and the emotion manifested is interpreted in an a shared way, by interactively building and linguistically labelling the schema that describes it. The input given to MIKE, on the contrary, is not a facial expression but a image that carries a 'visual impression' (a 'happy', a 'romantic', a 'cold', a 'fresh' image): this impression is defined, as well, in a shared way between system and User, through a set of visual characteristics of the image itself. Although the semantic repertoire of a image is much wider than face expressions, the two systems (MOUE and MIKE) are based on a common hypothesis: that languages employed to manifest affective states (metaphorical language, body language and verbal language) are ambiguous and may be misinterpreted and that this ambiguity may be reduced through an interactive negotiation between User and system. This is an interesting hypothesis indeed, that does not resolve, however, the question of whether the 'shared language' which emerges from this interaction will have a general value: i.e., whether this language will be User-independent or at least context-independent for the same User. This will certainly be verified when the two described systems will be applied to a number of cases larger than those that were examined in the study described in this paper.

This Special Issue ends with the reviews of three books: the first two, by Annika Waern, concern 'the' classic book in the field of affective interaction ('Affective Computing', by Rosalind Picard) and the more recent 'Human Cognition and Social Agent Technology', by Kerstin Dautenhan. The third review, by Kristina Hook, is about the last publication in the domain ('Affective Interactions', edited by Ana Paiva). The two books, together with the papers published in this Issue, offer us a number of reflections on what might be done to favour adaptation of human-machine interaction to affective factors: I wish to discuss briefly some of these open issues, in the hope that this may encourage new contributions from people in the User Modeling area. But, before this, a curious remark: what epidemiologists would call the 'sex ratio' of this Special Issue (and of contributions, in general,

to this particular domain) is strongly and evidently unbalanced in favour of women (with the very welcome exception of Cristoph Bartneck!). As Rosalind Picard says in her interview at the end of the book edited by Ana Paiva, *'In my own experience, I've never had such a high percentage of women in my research group as I've had since including affect on the agenda, but this is only one data point.'* Apparently, there are many more similar 'data points', all over the world.

2. Towards a Modeling of the *Cognitive Components* of Affective Factors

Although, as I said, we have recently seen a flourishing of research initiatives aimed at endowing machines with the ability to make 'affective' interactions with the User, research efforts in this domain are focused on 'communication aspects' and on the 'external expression' of affective factors rather than on their cognitive background and their behavioral consequences. Communication is one of the ingredients of personality: for instance, it is the main component of the 'extraversion' factor, in the FFM. However, traits and attitudes that may be relevant in HCI are not necessarily manifested through communication or through a clearly distinguishable external expression. For instance, a 'delegation' and 'help' trait or attitude is implicitly embedded into any User Interface, although often in a not conscious—and even not consistent—way (Castelfranchi et al., 1998). In addition, even if only communication aspects are considered, reacting appropriately to the Interlocutor's (human or artificial agent) manifestation of an affective state may require reasoning on the *causes* of this state rather than just recognising it. I have to understand the *reasons* of your fear to be effective in reassuring you; I have to understand exactly why you are angry and with which person if I want to appease you, and so on. Understanding this reason (and, often, even interpreting an expression) is, in general, not easy; in multimodal conversation or story understanding, interpretation of the subtle meaning of utterances is not immediate or unique. Ekman, for instance, showed that the same facial signal may occur in very different situations (Ekman & Friesen, 1978). Similar considerations can be done for the interpretation of sentences in natural language. Poggi and Pelachaud compared 'surprise' with 'doubt' and 'perplexity', to show that they share some features of the face expression (raising of eyebrows) and a common core of meaning, although they originate from different communication goals (Poggi & Pelachaud, 2000).

The only reasonable solution to this ambiguity problem is to try to come to a plausible interpretation by combining all information available. This means, first of all, combining verbal with nonverbal signs by giving an appropriate 'weight' to each of them, for instance with a belief network (as in Ball & Breese, 2000). However, in several situations, integrating information about these signs with knowledge about their possible meaning may contribute to further disambiguate the interpretation. This implies reasoning on the event, the action or the object (to employ OCC model) which probably produced the perceived state and on the mechanism which 'triggered' that state and the decision to manifest—or to

partially hide—it: the goals and beliefs that possibly produced it and their relation with the presumed mental state of the interlocutor. To this aim, some insight on *the cognitive background* of affective states is needed. Identical considerations may be forwarded for the generation of affective expressions or behaviors, that may take advantage of a representation of the situations in which they occur.

So, in the end, representing the cognitive aspects of affective states may ease and qualify their interpretation or generation. If the Speaker can reason on the Interlocutor's mind and try to infer what brought her to produce the (verbal and/or non-verbal) communicative act he received, he will have more chance of responding appropriately to the 'real' problem, rather than manifesting a generic 'empathy' towards that person. He will also reason on whether to manifest or not his emotion, according to its 'intensity', his personality, the context in which interaction takes place and so on.

We, in the User Modeling and Adaptation area, may find excellent hints for our work in this domain in the theories proposed by cognitive psychologists that I already cited in this Preface: the theory of personality and its relationship with goals by Carbonell (1980); the theory of 'subjective importance' of emotions by Ortony (1988); the emotion model architecture proposed by Frijda and Swagerman (1987); the theory of cognitive evaluation of social emotions by Castelfranchi (2000), and many more. Several useful references may be found in the papers published in this Issue and also in (Hudlicka & Fellows, 1996).

What we might do is to reconsider these theories and experiences in the light of the methods that have been applied successfully in User and context-adapted systems, during the last ten years. We will have to consider a number of questions, such as the following:

- Given that the main difference between representation of personality, mood (or attitude) and emotion seems to be related to their 'stability', may we conclude that this difference has something to do with long and short-term user modeling? if so, how may short and long-term features be integrated?
- How might distinct (either complementary or contradictory) affective aspects be assembled into a cognitively plausible combination within an Agent model?
- What are the advantages of the various methods that might be employed in building cognitive models of affective states and their relationship with non-affective components? For instance: are production rules, belief networks—static or dynamic—inference networks, neural networks etc suitable formalisms, and in which situations?

I do not pretend to provide an answer to these questions. I will just try to summarise some of the hypotheses forwarded by the above-mentioned theories, which should not be forgotten when making our methodological proposals, if we want to avoid the risk of oversimplifying the problem:

1. *Affective factors are linked to goals* (Carbonell, 1980; Castelfranchi, 2000);

- personalities may be defined in terms of *weights given to goal hierarchies*, of deviations from configurations of goals considered as ‘typical of a normal situation’;
- emotions *are triggered from* the belief about the *success/failure to achieve a goal* due to some ‘interrupt’ (change) in the external event or to some internal cognitive processing;
- emotions *control the process of goal achievement* by informing the individual whether the monitored goals are compromised or have been achieved;
- emotions *activate irrational goals* without the mediation of planning: positive appraisal activates ‘approach’ goals while negative appraisal activates ‘avoidance’ goals.

Notice that various characteristics of goals are important in the simulation of affective factors: having/not having a goal, with a given weight; goal activation/inactivation; goal achievement; goal monitoring.

2. *Time plays a crucial role in differentiating and characterising affective factors* (Pfeifer, 1988; Ortony, 1988; Lisetti, next issue; Castelfranchi, 2000):
 - the difference between personality traits, attitudes, moods and emotions may be described in terms of ‘time stability’.
 - the emotional state of an agent at a given time depends on his/her mental state at that time, on the events which occurred in the time immediately preceding it, but also on his/her previous emotional state; emotions evolve according to a decay/enhancement mechanism: ‘once generated, (they) do not stay around forever’.
3. *Context influences emotion triggering* (Ortony et al., 1988; Ortony, 1988; Elliott, 1994; Poggi & Pelachaud, 2000):
 - the external environment influences emotion elicitation through beliefs about events, objects, actions that may affect achievement or threatening of goals; these facts are evaluated in terms of their ‘desirability’, the degree to which they are realised at a given time, the degree of effort expended in attempting to prevent or to foster them etc;
 - the social environment influences emotion regulation. Personality and emotions have a *social valence*: that is, they are culturally praised and encouraged or criticised and discouraged. In addition, they are displayed in order to achieve some goal. As a consequence, their display is influenced by the personality and/or the affective state of the interlocutor(s), their role and the power relationships with them etc;
4. *Emotions may be activated endogenously*, not necessarily by external events: therefore, the monitoring mechanism that is responsible for deciding whether and how an emotion might be triggered should analyze the *internal mental state* of the agent, in addition to the external environment;

5. *Affective factors interact among themselves*: emotion triggering is influenced by the belief that achievement of a 'primary' goal (that is important for the Agent) is influenced (either directly or indirectly, through some 'instrumental' goal) by a given event; the weight given to primary goals is related, in its turn, to the personality. The two factors combine to define a general relationship between personality traits (or attitudes) and emotional states. In particular, one of the traits in the FFM (neuroticism, or 'emotional stability') contributes to defining a threshold in the activation of emotions, the propensity for a given emotional reaction;
6. *The combination of different (potential) affective states is a complex mechanism* (Carbonell, 1980): complementarity or contradiction between states may exist, and 'resolution criteria' that decide how to combine them have to be established;
7. *Emotions have an intensity* (Ortony, 1988; Elliott & Siegle, 1993): the agent's behavior depends not only on his/her *being in* some emotional state, but also in the intensity of this state: a 'threshold' effect regulates the decision to display it. An agent may be more or less ambitious, may feel more or less fear, etc. The face, gesture, voice manifestation of these states depends on its intensity, as well as on the plans or goals that are activated as a consequence of this state. Variables that are thought to affect the emotion intensity may be categorised according to their stability;
8. *Uncertainty cannot be ignored*: as the phenomenon is very complex (Pfeifer, 1985), pretending to model it in a purely logical way is probably not realistic. Uncertainty should be introduced in the 'triggering', the 'display' and the 'recognition' mechanisms.

Which formalisms may respond to these needs? Maybe different ones, according to the particular application domain and therefore to the particular aspect considered. Maybe hybrid models, as Pfeifer suggested might be useful. Here, the old and wellknown exhortation, by Elaine Rich, to '*be realistic*', is, once again, precious (Rich, 1988). In representing the so-called 'rational' aspects of the User, we have never tried to build complete models of their mind (what Borges would call 'a world map of natural size'): therefore, we probably don't need to model *all* possible affective states of the User. We should try to assess, first of all, whether affective factors *are really relevant* in the application domain considered. Then, ideally, we should accurately select those factors which have to be considered in that domain; for instance, fear or anxiety in doctor-patient interaction and, in general, in riskful situations. We should look for parameters that enable recognising that state *as much unobtrusively as possible*. Whenever possible, we should collect a corpus of data on how humans behave in similar circumstances and should analyse them to better understand the phenomenon. At this point, we would be ready to build our models and would, finally, *evaluate* the results obtained by comparing them with the corpus. This last point would require developing ad hoc methods, especially if the goal is to evaluate not only the 'convincingness' of an affective interface but also whether

affective factors do influence the Users behavior 'as if' they were interacting with a human, affective interlocutor.

Acknowledgement

The authors are grateful to Mrs Pauline Butts for her help in the preparation of the manuscript.

References

- André, E., Rist, T., Van Mulken, S., Klesen, M. and Baldes, S.: 2000, The automated design of believable dialogues for animated presentation teams. In: S. Prevost, J. Cassel, J. Sullivan and E. Churchill (eds.), *Embodied Conversational Characters*. MIT Press, pp. 220–255.
- Arafa, J., Charlton, P., Mamdani, A. and Fehin, P.: 1999, Designing and building Personal Service Assistants with personality. *Proceedings of the Workshop on Embodied Conversational Characters*, S. Prevost and E. Churchill (eds.), Tahoe City, pp. 95–104.
- Ardissono, L., Boella, G. and Lesmo, L.: 1999, Politeness and speech acts. *First Workshop on 'Attitudes, Personality and Emotions in User-Adapted Interactions'*, pp. 41–55.
- Ball, G. and Breese, J.: 2000, Emotion and personality in a conversational agent. In: S. Prevost, J. Cassel, J. Sullivan and E. Churchill (eds.), *Embodied Conversational Characters*. MIT Press, pp. 189–219.
- Bates, J.: 1994, The role of emotion in believable agents. *Communications of the ACM* **37**, 7, pp. 122–125.
- Berthold, A. and Jameson, A.: 1999, Interpreting symptoms of cognitive load in speech input. In: J. Kay (ed.), *Proceedings of User Modeling '99*, pp. 235–244.
- Carberry, S. and Schroeder, L.: 2001, Recognizing and conveying attitude and its underlying motivation. *Second Workshop on 'Attitudes, Personality and Emotions in User-Adapted Interaction'*.
- Carbonell, J. G.: 1980, Towards a process model of human personality traits. *Artificial Intelligence*, 15, pp. 49–74.
- Castelfranchi, C., de Rosis, F., Falcone, R. and Pizzutilo, S.: 1998, Personality traits and social attitudes in multiagent cooperation. *Applied Artificial Intelligence* **12**, 7–8, pp. 649–676.
- Castelfranchi, C.: 2000, Affective appraisal versus cognitive evaluation in social emotions and interactions. In: A. Paiva (ed.), *Affective Interactions*. Springer, pp. 76–106.
- Castelfranchi, C., de Rosis, F. and Carofiglio, V.: Can computers deliberately deceive? A simulation attempt for Turing's 'Imitation Game'. *submitted for publication*.
- Clore, G. and Ortony, A.: 1988, The semantics of the affective lexicon. In: V. Hamilton, G. H. Bower and N. H. Frjida (eds.), *Cognitive perspectives on emotion and motivation*. Kluwer Academic Publishers, pp. 367–394.
- de Rosis, F., Grasso, F. and Berry, D.: 1999, Refining instructional text generation after evaluation. *Artificial Intelligence in Medicine*, 17, pp. 1–36.
- de Rosis, F., Grasso, F., Castelfranchi, C. and Poggi, I.: 2000, Modeling conflict-resolution dialogs. In: R. Dieng and J. Mueller (eds.), *Computational Conflicts: Conflict modeling for Distributed Intelligent Systems*. Kluwer Academic Publishers, pp. 41–62.
- de Rosis, F. and Grasso, F.: 2000, Affective natural language generation. In: A. Paiva (ed.), *Affective Interactions*. Springer LNAI1814, pp. 204–218.
- Ekman, P. and Friesen, W.: 1978, *Unmasking the face: A guide to recognizing emotions from facial clues*. Prentice Hall Inc.

- Elliott, C. and Siegle, G.: 1993, Variables influencing the intensity of simulated affective states. *Proceedings of the AAAI Spring Symposium on Mental States '93*, pp. 58–67.
- Elliott, C.: 1994, Research problems in the use of a shallow Artificial Intelligence model of personality and emotion. *Proceedings of 12th AAAI Conference*, pp. 9–15.
- Elliott, C.: 1999, Why boys like motorcycles: using emotion theory to find structure in humorous stories. *Proceedings of the Workshop on EBAA '99*.
<http://condor.depaul.edu/~elliott/papers/humor/humor.html>
- Fogg, J. and Nass, C.: 1997, Silicon sycophants: the effects of computers that flatter. *International Journal of Human-Computer Studies*, 46, pp. 551–561.
- Frijda, N. H. and Swagerman, J.: 1987, Can computers feel? Theory and design of an emotional system. *Cognition and Emotion*, pp. 235–257.
- Hovy, E. H.: 1988, Affect in text. In: *Generating natural language under pragmatic constraints* (Chapter 4). Lawrence Erlbaum Associates, Hillsdale, NJ, pp. 54–80.
- Hudlicka, E. and Fellous, J.-M.: 1996, Review of Computational Models of Emotion. *Technical Report 9612*. Arlington, MA: Psychometrix Associates, Inc.
- Isbister, C. and Nass, C.: 1999, Personality in conversational characters: building better digital interaction partners using knowledge about human personality preferences and perceptions. In: S. Prevost and E. Churchill (eds.), *Proceedings of the Workshop on Embodied Conversational Characters*, pp. 105–114.
- Lester, J. C., Towns, S. G., Callaway, C. B., Voerman, J. L. and FitzGerald, P. J.: 2000, Deictic and emotive communication in animated pedagogical agents. In: S. Prevost, J. Cassel, J. Sullivan and E. Churchill (eds.), *Embodied Conversational Characters*. MIT Press, pp. 123–154.
- Loyall, A. B. and Bates, J.: 1997, Personality-rich believable agents that use language. *Proceedings of Autonomous Agents*, pp. 106–113.
- McCrae, R. and John, O. P.: 1992, An introduction to the Five-Factor model and its applications. *Journal of Personality* 60, pp. 175–215.
- Moon, Y. and Nass, C.: 1996, How ‘Real’ are Computer Personalities? *Communication Research* 23, 6, pp. 651–674.
- Moon, Y. and Nass, C.: 1998, Are computers scapegoats? Attributions of responsibility in human-computer interactions. *International Journal of Human-Computer Studies* 49, pp. 79–94.
- Nass, C., Moon, Y., Fogg, B. J., Reeves, B. and Dryer, C. D.: 1995, Can computer personalities be human personalities? *International Journal of Human-Computer Studies*. 43, pp. 223–239.
- Nass, C., Fogg, B. J. and Moon, Y.: 1996, Can computers be teammates? *International Journal of Human-Computer Studies*, pp. 669–678.
- Nass, C., Isbister, C. and Lee, E.-J.: 2000, Truth is beauty: Researching embodied conversational agents. In: S. Prevost, J. Cassel, J. Sullivan and E. Churchill (eds.), *Embodied Conversational Characters*. MIT Press, pp. 374–402.
- Ortony, A., Clore, G. L. and Collins, A.: 1998, *The cognitive structure of emotions*. Cambridge University Press.
- Ortony, A.: 1988, Subjective importance and computational models of emotions. In: V. Hamilton, G. H. Bower and N. H. Frijda (eds.), *Cognitive perspectives on emotion and motivation*. Kluwer Academic Publishers, pp. 321–343.
- Pfeifer, R. and Nicholas, D. W.: 1985, Toward computational models of emotion. In: L. Steels and J. A. Campbell (eds.), *Progress in Artificial Intelligence*. Ellis Horwood, pp. 184–192.
- Pfeifer, R.: 1988, Artificial Intelligence models of emotion. In: V. Hamilton, G. H. Bower and N. H. Frijda (eds.), *Cognitive perspectives on emotion and motivation*. Kluwer Academic Publishers, pp. 287–320.

- Poggi, I. and Pelachaud, C.: 2000, Emotional meaning and expression in animated faces. In: A. Paiva (ed.), *Affective Interactions*. Springer, pp. 182–195.
- Rich, E.: 1988, Stereotypes and User Modeling. In: A. Kobsa and W. Wahlster (eds.), *User Modeling in Dialog Systems*, pp. 35–40.
- Taylor, I. C., McInness, F. R., Love, S., Foster, J. C. and Jack, M. A.: 1999, Providing animated characters with designated personality profiles. *Proceedings of the Workshop on Embodied Conversational Characters*, S. Prevost and E. Churchill (eds.), Tahoe City, pp. 89–94.
- Turing, A. M.: 1950, Computing Machinery and Intelligence. *Mind*. London N.S.
- Walker, M. A., Cahn, J. A. and Whittaker, S. J.: 1997, Improvising linguistic style: Social and affective bases for Agent personality. *Proceedings of Autonomous Agents*, pp. 96–105.
- Wegman, C.: 1988, Emotion and argumentation in expression of opinion. In: V. Hamilton, G. H. Bower and N. H. Frjida (eds.), *Cognitive perspectives on emotion and motivation*. Kluwer, pp. 239–264.

Authors' vita

Fiorella de Rosis is a Full Professor of Informatics at the University of Bari; she coordinates the Intelligent Interfaces group and teaches at the Course of Human-Computer Interaction: Intelligent Interfaces, in the 5-year curriculum in computer science of that University.

After graduating in physics, with a dissertation in theoretical computer science, she worked for a few years at UNESCO and at the National Research Council in Rome, in the area of medical informatics.

Since 1980, she has been teaching at the University of Bari. Her research interests include uncertainty in user modeling, natural language generation and dialog simulation, with a particular interest towards 'believability' in interaction and the issues that make computer-generated messages or conversations similar to those occurring among humans (affect, deception, etc). She is the author of more than one hundred papers in the mentioned domains and is a member of the European Networks of Excellence AgentLink and i3net. More information may be found at the url: <http://aos2.uniba.it:8080/IntInt.html>