Developing a Consistent View on Emotion-oriented Computing

Marc Schröder[1] and Roddy Cowie[2]

[1] DFKI GmbH, Saarbrücken, Germany
[2] Queen's University, Belfast, Northern Ireland
http://emotion-research.net

Abstract. The network of excellence HUMAINE is currently making a co-ordinated, interdisciplinary effort to develop a consistent view on emotion-oriented computing. This overview paper proposes a "map" of the research area, distinguishing core technologies from application-oriented and psychologically oriented work. First results from the on-going research in the thematic workpackages are reported.

1 Introduction

It is increasingly recognised that emotional factors in a broad sense are central to improving the naturalness of interaction between machines and their users. As humans, users react emotionally to aspects of their environment that matter to them [1], and these emotions will influence their way of acting [2], their way of thinking [3], and their decisions [4]. Furthermore, as social animals, humans expect their interaction partners to pick up signs of emotion and to react to them in some appropriate way [5]. Currently machines that interact with human users do not take account of the emotional dimension that humans expect to find in interaction, and that is a recurrent source of frustration. A simple example for emotionally inadequate system behaviour is a message window suggesting a software update which is triggered while the user is under time pressure or giving a presentation: it is likely to induce panic or anger rather than appreciation. A more sophisticated example is a multi-modal dialogue system that cannot anticipate the emotional impact of a piece of information on the user: if the system informs the user that, e.g., there are no more seats left on the flight the user wants to book, a standard happy-sounding voice will not improve customer relations.

Creating competent emotion-oriented systems is a large scale challenge. The European Network of Excellence HUMAINE (HUman-MAchine Interaction Network on Emotions) was established to prepare the scientific and technological ground for this task, with funding from the EU IST programme from 2004 to 2007.

HUMAINE follows a principled approach to addressing the large number of issues involved. As the network brings together researchers with a very wide range of backgrounds, a first phase of 18 months was scheduled in the work plan,

whose aim was to identify an appropriate set of sub areas into which research can be structured, and to come to a common understanding of the core issues in each of these thematic areas. In several iterations, this first phase led to the establishment of plans for "exemplars", i.e. achievable pieces of work illustrating how things should be done in a principled way in a given area. In each thematic area, a workshop was held (or will soon be held), in which the subject matter is highlighted from a broad range of perspectives. The proceedings of these workshops are available on the HUMAINE portal (`http://emotion-research.net/ws`). The discussions of phase one are now basically complete, and the second phase is starting, in which these exemplars are actually built.

This paper reports on some of the key outcomes of the first phase, mainly by outlining the "exemplars" envisaged in order to advance the state of the art in the different thematic areas.

2 Mapping the Research Area

A key part of the first project phase is reaching an understanding of the various tasks and disciplines relevant to emotion-oriented computing, and the ways in which they may interact. The ideas incorporated in the initial HUMAINE proposal have fared reasonably well, but they have developed with experience. Figure 1 summarises what seem to be the key divisions and connections at this stage.

The central column represents the areas where purely technological challenges loom largest. It is not self-evident that detection and synthesis should function as separate sub-disciplines, and HUMAINE initially proposed a different division: but it has become clear that they draw on different background technologies. 'Planning action' involves modelling the kind of action pattern that might be expected in a particular emotional state, either so that an artificial agent can generate appropriate action patterns or so that it can anticipate the kinds of action pattern that a human might produce in a given state - which in turn may be used to recognise emotion or to select among various responses that might be considered.

The left hand column deals with issues where application is most obviously of concern. The special character of usability issues in this area has gradually become clearer. Finding out how users respond to an emotion-oriented system is both more difficult than it is for technologies with more objective aims, and more important. It is more difficult because emotional responses are subtle, and easily disrupted by interventions that are meant to measure them. It is more important because designing an emotion-oriented system is centrally concerned with accommodating to non-rational preferences and dispositions in the user. In that situation, iterative user-centred design methods seem likely to be indispensable. Work on emotion-related language could in principle be divided up among the sub-disciplines in the central column, but in practise it draws on different conceptual roots and has particular links to applications in the relatively near

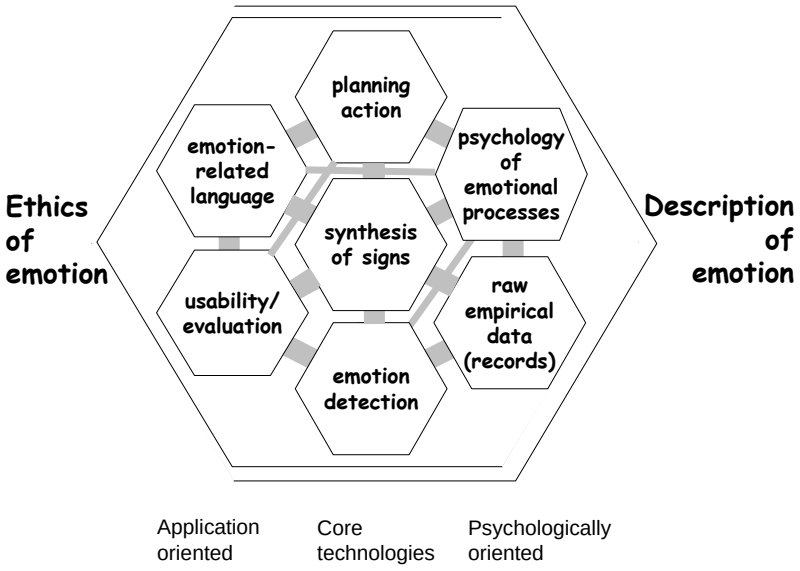

Fig. 1. Graphical representation of the sub-areas involved in achieving emotion-oriented computing

future. It has become clear that HUMAINE's workpackage on persuasion and communication in effect represents that area.

The right hand column contains the sub-areas with the strongest roots in psychology. These are divided between one concerned with theory and one concerned with empirical data, which was described in the original proposal as 'databases'. Of course existing theory is informed by data, but for a variety of reasons, it has not generally been the kind of data that it is natural to collect with a view to developing emotion-oriented systems. As a result, there are creative tensions between that kind of data collection and existing psychological theory. Similarly, psychological theory does not simply inform technological work. Technological work is a test of its accuracy and completeness in a general sense. It also promises to provide unparallelled tools to test theories, in the form of artificial agents whose actions can be controlled with a precision that is impossible with humans, and which therefore provide unparallelled opportunities to test theoretical ideas.

It is not an accident that synthesising signs of emotion is at the centre of the diagram. There are rich connections among all the areas, but the task of synthesising agents that can interact emotionally seems at the present time to

be the one that best summarises the state of the art - in the sense that it cannot be done well without satisfactory progress in all the others. Conversely, failures in this area may expose problems in any of the other areas.

At either edge of the diagram are issues with a strong philosophical element which affect the whole enterprise, though they do not impact all of the areas equally. One, whose strongest effect is on the areas related to psychology, is finding appropriate ways to describe emotions and emotion-related states. The other, whose strongest effect is on the areas close to application, is the ethics of emotion-oriented systems.

This summary is deliberately at a broad level. The sections that follow look at individual thematic areas in more detail.

3 Describing Emotions

It is difficult to work with phenomena unless one has good ways of describing them. In the broad area of emotion, it is widely accepted that words in common use are an unsatisfactory medium. It is not simply that potentially interesting states elude description: even core terms like "emotion" or "affect" are defined differently by different experts, and they point non-experts in various directions according to context [6, 7]. HUMAINE has chosen to confront the issue by making the description of emotion an explicit focus of research. It is being addressed on several levels.

At the first level, there has been vigorous discussion within HUMAINE about the domain that needs to be addressed [8]. The surface argument involves questions about how widely or narrowly to use words like 'emotion' and 'affect', but it reflects a deeper issue, which is to mark out a practically important set of problems that yield to conceptually related solutions. As a first step, HUMAINE aims to map out states which are or might be considered to be related to emotion, and which are both reasonably common and potentially relevant in human-computer interaction. HUMAINE aims to carry out that kind of 'actuarial' exercise, and initial steps have been taken [9]. Words then need to be found (or invented) that allow the domains that matter practically to be described in a way that is convenient and intuitive.

At a second level, the Theory workpackage in HUMAINE has started to build up a glossary of emotion-related terms, drawing on the existing literature [10]. This conceptual and terminological clarification task will differentiate the different types of affective phenomena that should be distinguished from a theoretical perspective, while at the same time staying close to application concerns. Comments from the potential users of this vocabulary (e.g., the "system-builders" in the network) will be sought in several iterations, leading to successive clarification and re-formulation, culminating in definitions that can be used by researchers from any theoretical background. Establishing a common terminology will significantly reduce the barrier to cross-disciplinary cooperation in this area.

The third level involves the practical methods of describing individual emotional states. Newcomers to the field tend to use short lists of basic emotions,

often not realising that these were proposed in the specific context of evolutionary, "Darwinian" emotion theories [6]. A key aim of the Theory workpackage is to promote awareness of the available choices of emotion descriptions, and the circumstances in which they may be useful. For example, lists of emotion words are being established that are particularly useful in the context of emotion-oriented technological systems [8]; descriptions of individual cognitive appraisal components can be related to aspects of facial expression; and broad dimensional labelling can be used to track the general emotional tone over time. As these options become better known in the technological community, system-builders can select the most suitable formalisms for their application.

The Databases workpackage in HUMAINE approaches the issue of emotion descriptions from a different angle [11]. Rather than formulating and specifying descriptions out of theory, this workpackage will explore the phenomena that can be found in "real" data, i.e. in naturalistic recordings of persons experiencing emotions [12]. Often, the phenomena observed in such contexts differ substantially from the clear-cut ideas that come to the fore in emotion theories. For example, emotion-related states of low intensity, as they often occur in natural dialogue, cannot easily be described by emotion theories which focus on intense, fullblown emotions. In addition, word-based emotion representations ("anger", "sadness" etc.) cannot easily capture the composites and shades of emotion that are frequently observed in naturalistic data, e.g. when two emotions are simultaneously present, or when one emotion is expressed in order to mask another one that is experienced [13]. One of the aims of the Databases workpackage is to provide "provocative" data which exposes these issues. On that basis, it aims to identify labelling schemes that are genuinely suited to describe that kind of material. These are expected to combine elements of several existing approaches – description using everyday verbal categories, broad dimensional descriptions, and descriptions based on appraisal theory.

4 Detecting and Generating Signs of Emotion

One fundamental of natural human-machine interaction is the ability to detect the signs of emotion emitted by the user, intentionally or unintentionally. This task is, on the one hand, heavily dependent on the emotion models used, e.g. whether the emotion is described as a category or as a region in a multi-dimensional space. The task is also highly dependent on the material from which to recognise emotion: for example, classifiers that work very well with acted emotional material may fail on naturalistic material [14]. The challenge involves events with an emotional component as well as "pure"emotions, such as dialogue success [15], and classification may well depend on contextual knowledge as well as local features [16].

HUMAINE has prioritised three modalities for study – facial expression in video; speech in audio; and physiological parameters. In each of these, its first priority is to establish the reliability of alternate signal analysis algorithms for extracting basic features. Building on that, it aims to clarify the principles of ef-

fective cross-modal integration. One challenge that is immediately apparent is to deal with the different temporal structures that characterise the modalities. Most algorithms in biosignal processing operate on a continuous, and relatively slow time scale. Speech tends to be analysed in discrete units, ranging from phonemes to phrases and even sequences of phrases; whereas the best known approaches to facial expression analysis deal with essentially instantaneous 'stills'. Beyond that, several theoretical models of multimodal integration will be compared. The simplest types of integration model to be considered include the "direct identification" model, where all input signals are directly transmitted to one multimodal classifier, and the "separate identification" model, where emotion is recognised from each modality separately. More complex integration models include the "dominant modality recoding" model, where a dominant modality drives the perception of other modalities, and the "common space recoding" model, where all the modalities are projected upon a common space prior to categorisation (as audiovisual speech recognition is thought to involve mapping information from both modalities into a common motor space).

Cutting across these issues is the question of which type of emotion descriptor to predict. Key options include discrete, dimensional and appraisal models of emotions. There is evidence that some types of descriptor relate particularly directly to the information available in some modalities [17]. HUMAINE will study the extent to which there are privileged relationships between descriptors and particular sources of evidence.

There is no unambiguous measure of success for such emotion recognition components: Should it perform as closely as possible like a human, i.e., make the same errors as humans, or should it be as accurate as possible, i.e. possibly more accurate than a human? The answer to this question is likely to depend on the application area. In a stress detection module for drivers, high accuracy is important; in a conversational interface, acting as human-like as possible will be more important. As a baseline for comparison, work is being performed within the Theory workpackage to assess the recognition capabilities of humans. Presenting the same audiovisual material to both human raters and classification algorithms will provide interesting insights in the similarity of their judgements.

Synthesising emotional signs is as important as their recognition for natural human-machine interaction. From a human point of view, it could be seen as just "the other side of the medal". However, the technologies involved in the two endeavours differ completely. HUMAINE investigates how emotions can be expressed by Embodied Conversational Agent (ECA) systems [18]. This work has started by compiling a list of capabilities that go beyond the current state of the art, but which would be required for an emotionally competent ECA. They can be grouped into three areas: perception, interaction, and generation. In the perception domain, an important pre-requisite for believable emotional interaction is an ECA's capability to perceive the user, events, or other agents. A key means for modelling this capability is an affect-related attention mechanism. On the level of interaction, rather than modelling the ECA merely as a speaker, it is important to attempt the generation of *listener* behaviour. Among other

things, this includes backchannel utterances, which can signal to the user that the ECA is listening; what it does or does not understand; and how it evaluates what is being said [19]. On the generation side, the existing capabilities such as gestural and vocal expressivity need to be refined both in richness and in control, in order to model more closely what humans do in expressive situations.

This last point is obviously linked to the analysis of naturalistic databases of emotional behaviour. Conceptually, the insights gained from such analyses can be used for improving the rendered ECA behaviour. But the link can be made even more direct: given a suitable description format, a database annotation can be used to "drive" an ECA, so that the ECA can serve as an analysis-by-synthesis framework for validating the annotation scheme. First promising steps in this direction have been undertaken [20].

5 Emotions in Computational Cognitive Architectures

Going beyond shallow descriptions of emotions requires an understanding of the emotional aspects of the cognitive architecture that processes them. Existing systems that predict emotional reactions from situation descriptions are often based on the cognitive emotion model proposed by Ortony, Clore and Collins [21], which provides a useful but limited account.

HUMAINE attempts to understand and describe the emotional aspects of cognitive architectures more fully by a combination of two approaches.

Conceptually, a "blueprint" description for an affectively competent agent will be compiled in the Theory workpackage [10], with the aim to describe the mechanisms involved in emotional processing. It will gather the different points of view taken by current theorists from different disciplines, including psychology, cognitive neuroscience, philosophy, and ethology, and should be the starting point for a fruitful dialogue with engineer-oriented groups regarding the issues encountered during the implementation process. It should be seen as an evolving source "book", where state-of-the-art questions could be asked and where attempts to address them will be described.

Practically, HUMAINE also works with existing cognitive system architectures, and investigates the various ways in which emotions can be incorporated in such architectures [22]. The workpackage on emotion in Cognition and Action explores emotional phenomena in a range of very different approaches to cognitive architectures.

The "low-level" or sub-symbolic approach is concerned with the investigation of the influence of emotions in cognition and action from the perspective of their embodiment. Following this view, cognition and action are inseparable, tightly coupled perception-action loops rather than separable input-output elements. In this line of research, HUMAINE investigates robots endowed with a relatively simple cognitive architecture inspired by biological perception-action loops. First results indicate that minor modifications to the architecture, modelled after biological neuromodulation, can give rise to emergent "emotional" behaviour. Ethological methods, which are usually applied to studying animal

behaviour, can be used to interpret the behaviour of such a robot. For example, in one scenario, two motivated robots competed for a resource ("food"). Usually, their bumpers would signal an obstacle that must be avoided. However, when the architecture was altered so that the bumper sensitivity dropped when the "hunger" became too big, the robots could be observed to show "aggressive" behaviour when trying to attain the resource, pushing each other away [23]. Another example is a visual homeostasis mechanism leading to "bonding" behaviour as displayed by Lorentz' geese [24].

The "high-level" or symbolic approach works on extensions of existing belief-desire-intention (BDI) models with emotion. This approach is studied in HUMAINE in the scope of dialogue simulation, and is based on Bayesian networks as a method to represent uncertain knowledge and reasoning. Existing cognitive models of emotion activation are being revised in the light of a document [25] produced by the Theories workpackage concerning emotion theories and classification of emotion models. Furthermore, it is envisaged to design an affective user modelling component to be combined with a linguistic parser of the user moves, in order to integrate 'recognition' and 'interpretation' functions. How to integrate these in a single cognitive model is another research question that will be addressed.

The "hybrid" approach attempts to bridge the gap between sub-symbolic and symbolic aspects of cognition and emotion. It will do so using agents that combine both levels, being at the same time embodied/reactive and deliberative. Studying the way in which these two levels are interrelated with respect to emotions first requires clear definitions of key properties, such as representational features, degree of autonomy, or independence from the outside world. On that basis, one can then start to deal with issues such as management of timing and prioritisation between both levels, or how to achieve mappings between the contents of these levels. For example, it would be interesting to know that an avoidance tendency on the reactive level is somehow linked to an emotion named "fear" on the symbolic level. Making this link explicit is not trivial. Similarly, where the reactive component generates a behaviour that appears contradictory to the symbolic emotional assessment of the situation, it will be most relevant to model the negotiation and decision-making process required between the two levels to generate behaviour.

All these approaches focus on modelling the cognitive mechanisms and resulting behaviour of an individual. Complementarily, HUMAINE also addresses the social and interpersonal mechanisms of regulating the emotions of individuals. Here, the unit of analysis is the *relation* in groups of two or more agents rather than the behaviour or mental state of an individual. This work includes a broad range of aspects, from the study of human politeness [26] to modelling the links between personality, emotion, and mood [27].

6 Affecting the User

The ultimate goal of HUMAINE is to enable the community to build emotion-oriented technological systems. Even if this is still an ambitious, long-term goal, work is underway in HUMAINE to prepare the ground.

The workpackage Emotion in Communication and Persuasion explores ways to purposefully induce emotions in the human user [28, 29]. Models of persuasion are developed and tested in both monological and dialogical situations. Natural language is one important means for inducing emotions in human users. A first application in this area is the creative humour testbed, where the system generates potentially funny slogans based on semantic properties of natural language [30]. This example shows the direction in which first applications of emotion-oriented technology are emerging: Rather than fully competent stand-alone systems, the language-based creative humour testbed can produce a number of *potentially* funny slogans out of which the really funny ones need to be selected by a human user.

Several types of communicative strategy are investigated in the context of persuasion. One such strategy is politeness, currently being explored in a conversational context [26]. Another is deception: the recognition and generation of lying expressions. As a scenario for investigating lying behaviour in human-machine interaction, an interactive dice game was developed, played by two humans and an ECA. The game can only be won by lying occasionally [31]. In this controlled scenario, various aspects of lying can be studied, related to system behaviour (simulated lying), user reactions to simulated lying, and user lying.

Affective issues in user interfaces present a new set of challenges to usability research. Work in the Usability workpackage focuses on finding methods that can help guide future design and evaluation of affective systems [32, 33]. Existing usability criteria such as control, predictability or transparency are not the most suitable for describing emotional systems. For that reason, the exemplar in this workpackage will first of all develop a set of criteria by which to measure successful, usable affective interaction systems. These criteria will not be objective, independently measurable entities, but will make sense relative to the specific application domain, aim to capture subjective experiences of the user, and foremost, be related to the designer's intention for the application. These criteria will then need to be translated into evaluation metrics, accompanied by suggested evaluation methods. Existing user-centred methods for design and evaluation will be investigated with respect to their use for emotion-oriented systems. At the same time, new methods will be proposed that are targeted specifically at capturing the unique aspects of affective interaction. Examples for such methods are: a sensual method for non-verbal mediation of affective state, a Wizard-of-Oz environment for multimodal emotional interaction, and an extended think-aloud protocol designed to capture emotional interactions.

7 HUMAINE Conscience

When dealing with machines that might one day be able to influence human emotions, there is a real need to think about the ethical dimension of such systems. It is only too easy to imagine, e.g., persuading machines used for "enhancing" product sales, or surveillance systems measuring continually the degree of friendliness exhibited by call-centre staff.

HUMAINE takes a proactive approach to these issues in its Ethics workpackage – few projects in the IST domain investigate ethical implications so thoroughly. An ethical audit [34] marked the starting point of this endeavour: it consisted of a questionnaire assessing participants' previous experience with ethical issues in emotion-oriented systems and in emotion research including human participants. It was completed by all HUMAINE partner institutions. Its results showed a serious lack of preparation among the organisations carrying out research in the area. More importantly, the parties involved were not necessarily aware of this shortcoming.

The challenge faced by HUMAINE's Ethics team is thus to set up a whole new body of procedures and criteria by which to make sure that research and its results are not used to put humans to unethical risks. The theoretical framework in which these issues are now being addressed [35] is called Principlism [36]. It is based on the four universally shared moral principles of nonmaleficence, autonomy, beneficence, and justice. It is acknowledged that applying these principles to a concrete situation involves delicate judgments. Any set of general recommendations will therefore need to be complemented by a panel of humans, e.g. by an ethics committee that can address specific situations.

8 Conclusion: Steps Ahead

This paper has given a short overview of the broad range of activities under way in the Network of Excellence HUMAINE. Despite the multitude of angles from which network members address the complex set of thematic areas, people have come to a common understanding of key research issues. Joint specification of plans for exemplars has been a crucial mechanism to bring perspectives closer.

In the second phase of HUMAINE, which is starting now, these exemplars will actually be built. Due to its nature as a network, HUMAINE will not produce full-scale demonstrator or prototype systems. This allows us to avoid the need to make the usual short-term shortcuts required to make a system look coherent. Instead, we will produce illustrations at various levels of technological sophistication. The core intention behind building this type of exemplars is to do things "right", in the interest of a well-founded, iterative build-up of competences. With this approach, we believe we can make a real, lasting contribution.

Acknowledgements

The preparation of this paper was supported by the EU project HUMAINE (IST-507422).

References

1. Scherer, K.R.: Introduction: Cognitive components of emotion. In Davidson, R.J., Goldsmith, H., Scherer, K.R., eds.: Handbook of the Affective Sciences. Oxford University Press, New York, Oxford (2003) 563–571
2. Frijda, N.H.: The Emotions. Cambridge University Press, Cambridge, UK (1986)
3. Sokolowski, K.: Emotion und Volition. Hogrefe, Göttingen (1993)
4. Damasio, A.R.: Descartes' Error: Emotion, Reason, and the Human Brain. Grosset/Putnam, New York (1994)
5. Frijda, N.H., Mesquita, B.: The social roles and functions of emotions. In Kitayama, S., Markus, H.R., eds.: Emotion and Culture. United Book Press, Baltimore, MD, USA (1994) 51–87
6. Cornelius, R.R.: The Science of Emotion. Research and Tradition in the Psychology of Emotion. Prentice-Hall, Upper Saddle River, NJ (1996)
7. Russell, J., Barrett-Feldman, L.: Core affect, prototypical emotional episodes and other things called emotion: dissecting the elephant. Journal of Personality and Social Psychology **76** (1999) 37–63
8. Cowie, R., Douglas-Cowie, E., Cox, C.: Beyond emotion archetypes: databases for emotion modelling using neural networks. Neural Networks (Special Issue Emotion and Brain: Understanding Emotions and Modelling their Recognition) (2005) to appear.
9. Douglas-Cowie, E., Cowie, R.: Report on the first HUMAINE summer school. http://emotion-research.net/ws/summerschool1 (2004)
10. Scherer, K.R., Roesch, E., et al.: HUMAINE deliverable D3d: Description of potential exemplars: "WP3: Theories and Models of Emotion". http://emotion-research.net/deliverables (2004)
11. Douglas-Cowie, E., et al.: HUMAINE deliverable D5d: Multimodal data in action and interaction: a library of recordings and labelling schemes. http://emotion-research.net/deliverables (2004)
12. Douglas-Cowie, E., Campbell, N., Cowie, R., Roach, P.: Emotional speech: Towards a new generation of databases. Speech Communication Special Issue Speech and Emotion **40** (2003) 33–60
13. Abrilian, S., Devillers, L., Buisine, S., Martin, J.C.: EmoTV1: Annotation of real-life emotions for the specifications of multimodal affective interfaces. In: Proceedings of HCI International 2005, Las Vegas, USA (2005) to appear.
14. Batliner, A., Fischer, K., Huber, R., Spilker, J., Nöth, E.: Desperately seeking emotions or: Actors, wizards, and human beings. In: Proceedings of the ISCA Workshop on Speech and Emotion, Northern Ireland (2000) 195–200
15. Steidl, S., Hacker, C., Ruff, C., Batliner, A., Nöth, E., Haas, J.: Looking at the last two turns, I'd say this dialogue is doomed - measuring dialogue success. In P. Sojka, I.K., Pala, K., eds.: Proceedings of Text, Speech and Dialogue, 5th International Conference, September 9-12, 2002, Brno, Czech Republic. Lecture Notes in Artificial Intelligence, Berlin, Springer (2004) 629–636
16. Wallace, M., Raouzaiou, A., Tsapatsoulis, N., Kollias, S.: Facial expression classification based on MPEG-4 FAPs: The use of evidence and prior knowledge for uncertainty removal. In: Proceedings of the IEEE International Conference on Fuzzy Systems (FUZZ-IEEE), Budapest, Hungary (2004)
17. Laukka, P.: Vocal expression of emotion: discrete-emotions and dimensional accounts. PhD thesis, Uppsala University (2004)

18. Pélachaud, C., et al.: HUMAINE deliverable D6c: Description of potential exemplars: Interaction. http://emotion-research.net/deliverables (2004)
19. Pfleger, N., Alexandersson, J.: Modeling non-verbal behavior in multimodal conversational systems. Information Technology **46** (2004) 341–345
20. Pélachaud, C.: Non-verbal communication. Presentation at the workshop of the HUMAINE workpackage 6, Emotion in Interaction, Paris, France (2005) http://emotion-research.net/ws/wp6.
21. Ortony, A., Clore, G.L., Collins, A.: The Cognitive Structure of Emotion. Cambridge University Press, Cambridge, UK (1988)
22. Cañamero, L., et al.: HUMAINE deliverable D7c: Description of potential exemplars: Emotion in Cognition and Action. http://emotion-research.net/deliverables (2004)
23. Avila-Garcia, O., Cañamero, L.: Using hormonal feedback to modulate action selection in a competitive scenario. In: From Animals to Animats: Proceedings of the 8th International Conference of Adaptive Behavior (SAB'04), Cambridge, MA, MIT Press (2004) 243–252
24. Blanchard, A.J., Cañamero, L.: Using visual velocity detection to achieve synchronization in imitation. In: Proceedings of the 3rd International Symposium on Imitation in Animals and Artifacts, Symposium of the AISB'05 Convention, Hertfordshire, UK (2005)
25. Scherer, K.R., et al.: HUMAINE deliverable D3c: Preliminary plans for exemplars: Theory. http://emotion-research.net/deliverables (2004)
26. Rehm, M., André, E.: Informing the design of embodied conversational agents by analysing multimodal politeness behaviors in human-human communication. In: Proc. AISB 2005 Symposium on Conversational Informatics for Supporting Social Intelligence & Interaction. (2005)
27. Gebhard, P.: ALMA – a layered model of affect. In: Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-05), Utrecht (2005) To appear.
28. Stock, O., et al.: HUMAINE deliverable D8a: Report on basic cues and open research topics in Communication and Emotions. http://emotion-research.net/deliverables (2004)
29. Stock, O., et al.: HUMAINE deliverable D8c: Description of potential exemplars: Communication & Persuasion. http://emotion-research.net/deliverables (2004)
30. Stock, O., Strapparava, C.: The act of creating humorous acronyms. Applied Artificial Intelligence **19** (2005) 137–151
31. Rehm, M.: Keeping users engaged – analysing gaze behaviours in a multiparty interaction with an eca. Presentation at the workshop of the HUMAINE workpackage 6, Emotion in Interaction, Paris, France (2005) http://emotion-research.net/ws/wp6.
32. Höök, K., et al.: HUMAINE deliverable D9c: Description of potential exemplars: Usability. http://emotion-research.net/deliverables (2004)
33. Höök, K., et al.: HUMAINE deliverable D9a: Usability of emotion oriented systems recommendations green paper. http://emotion-research.net/deliverables (2005)
34. Goldie, P., Döring, S.A.: HUMAINE deliverable D0f: Report on Ethical Audit. http://emotion-research.net/deliverables (2004)
35. Goldie, P., Döring, S.A., et al.: HUMAINE deliverable D10b: Interim report to plenary meeting on ethical frameworks for emotion-oriented systems 2005. http://emotion-research.net/deliverables (2005)
36. Beauchamp, T.L., Childress, J.F.: Principles of Biomedical Ethics. 5th edn. Oxford University Press, Oxford (2001)