Expressive Agents: Non-verbal Communication in Collaborative Virtual Environments

Fabri M*, Moore DJ*, Hobbs DJ[†]

*Leeds Metropolitan University, [†]University of Bradford

1. ABSTRACT

The premise of this paper is that agent technology in collaborative virtual environments (CVEs) may be enriched by incorporating an emotional channel alongside the conventional informational content, and that this would be best achieved through an associated visual human embodiment or avatar. Since humans express emotion in face-to-face encounters primarily through facial expression, an investigation was undertaken in order to establish how such expressions might be effectively and efficiently captured and represented visually. The study involved consulting socio-psychological research relating to face-to-face encounters, followed by an experimental study to establish user ability to interpret the faces of the avatars pre-prepared to express specific emotions. Effectiveness was demonstrated through good recognition rates for all but one of the emotion categories, and efficiency was established since a reduced feature set was found to be sufficient to build the successfully recognised core set of avatar facial expressions.

2. PROBLEM CONTEXT

Current forms of electronic communication generally lose the emotional context, along with the ability to express emotional states in the ways experienced in face-to-face conversations. Text-based tools are notoriously unreliable at conveying emotion (Ødegård 1993, Lisetti et al 2001). Audio conferencing tools can ease some of these difficulties but lack ways of mediating non-verbal forms of communication. On the other hand, in face-to-face interactions, facial expression, posture and gesture play an important and significant role. Such *socio-emotional content* (Lisetti et al 2001) is vital for building trusting, productive relationships that go beyond the purely factual and task-oriented communication.

Indeed, Morris et al (1979) hold the view that these non-verbal signals are even more important than verbal information, particularly in respect of the communicating of changing moods and emotional states. Social psychologists assert that more than 65% of the information exchanged during a person-to-person conversation takes place in the non-verbal band (Knapp 1978, Morris et al 1979). In a learning environment, for example, the ability to show emotion, empathy and understanding through facial expressions and body language is central to ensuring the quality of tutor-learner interaction (Cooper et al 2000). Recent findings in psychology and neurology suggest that emotions are also an important factor in decision-making, problem solving, cognition and intelligence in general (Picard 1997, Lisetti and Schiano 2000, Damásio 1994, Dittrich 1993).

Collaborative Virtual Environments (CVEs) aim to reintroduce emotional and social context to distance communication whilst at the same time offering a stimulating and integrated framework for conversation and collaboration. Users can become actively engaged in interaction with the virtual world and with other inhabitants. Again, in a distance learning discipline, this high-level interactivity with the users' senses, is seen as an essential factor for effective and efficient learning (Stoney and Wild 1998).

2.1. NON-VERBAL COMMUNICATION IN SOCIAL INTERACTION

The term "non-verbal" is commonly used to describe all human communication events that transcend the spoken or written word (Knapp 1978). Argyle (1988) sees non-verbal behaviour taking place primarily through facial expression, bodily contact, gaze (and pupil dilation), spatial behaviour gesture, clothing and appearance, body posture, and non-verbal vocalisation. When two parties interact, they monitor and interpret each other's emotional expression (Strongman 1996); hundreds of expressive movements are employed every day as part of the social interaction events of a typical day (Morris et al 1979), and their correct use is an essential part of our social competence and skills.

Non-verbal communication typically serves to repeat, contradict, substitute, complement, accent, or regulate verbal communication (Knapp 1978). Kendon (1983) argues that employing non-verbal means of expression is not even necessarily dependent upon the use of verbal language - non-verbal communication can be separate from, and in principle equal to or more effective than speech. Dittrich et al (1996) consider the ability to judge the emotional state of others an important goal in human perception, and it is stressed that from an evolutionary point of view, it is probably the most significant function of interpersonal perception. Since different emotional states are likely to lead to different courses of action, it can be crucial for survival to be able to recognise emotional states, in particular anger or fear, in another person.

All the different channels for non-verbal communication - face, gaze, gesture, posture - can in principle be mediated in CVEs. For the current experimental study, however, the focus was on the face since in the real world it is this channel that is the most immediate indicator for the emotional state of a person (Ekman and Friesen 1975). The face reflects interpersonal attitudes, provides feedback on the comments of others, and is regarded as the primary source of information after human speech. For these reasons, humans naturally pay a great deal of attention to the messages they receive from the faces of others (Knapp 1978).

On the basis of this social science research, it is argued that this naturally developed skill to "read" facial expressions is potentially highly beneficial to communication in CVEs. It is further proposed that the emotionally expressive virtual face on the avatar of an interlocutor, whether human or agent-based, may be able to aid the communication process and provide information that would otherwise be difficult to mediate.

2.2. THE NEED FOR EXPRESSIVE AGENTS

The developing model of CVEs includes intelligent aspects of a user provided through a suitably programmed agent, and visual representation of the user embodiment through an avatar. Such visual representations, however, remain relatively simple and rudimentary (Thalmann 2001). In particular, virtual environments are often poor in terms of the emotional cues that they convey (Fleming and Dobbs 1999). Accordingly, the need for sophisticated ways to reflect emotions in virtual embodiments has been pointed out repeatedly in recent investigations (Dumas et al 1998, Guye-Vuillème et al 1999).

McGrath and Prinz (2001) call for appropriate ways to express presence and awareness in CVEs in order to aid communication between inhabitants, be it full verbal communication or non-verbal *presence in silence*. Thalmann (2001) sees a direct relation between the quality of a user's representation and their ability to interact with the environment and with each other. Durlach and Slater (1998) observe that even avatars with rather primitive expressive abilities may engender strong emotional responses in people using a CVE system. It appears, therefore, that the avatar can readily take on a personal role, thereby increasing the sense of togetherness or community feeling. It potentially becomes a genuine representation of the underlying individual, not only visually, but also within a social context.

However, this does not necessarily imply that a 'good' avatar has to be a photorealistic and accurate representation of human physiognomy. There has been early evidence that approaches aiming to reproduce the human physics in detail may in fact be wasteful and counterproductive (Benford et al 1995). Hindmarsh et al (2001) suggest that a straightforward translation of human physical embodiments in CVEs is likely to be unsuccessful, at least until full perceptual capabilities of physical human bodies are also available in virtual space. Even then, opportunities for employing more inventive and evocative modes of expression would probably be lost if the focus is merely on simulating the real world.

It may be more appropriate, and indeed more supportive to perception and cognition, to represent users in simple or unusual ways. Godenschweger et al (1997) found that minimalist drawings of body parts, showing gestures, were generally easier to recognise than more complex representations. Donath (2001) warns that because the face is so highly expressive and we are so adept in reading (into) it, any level of detail in 3D facial rendering could potentially provoke the interpretation of various social messages. If these messages are unintentional, the face is arguably hindering communication more than it is helping. Also, there is evidence that exaggerated or particularly distinctive faces can convey emotions more efficiently than normal faces (Bartneck 2001, Zebrowitz 1997, Ellis 1990), a detail regularly taken advantage of by caricaturists.

The aim of this research project was thus to investigate the use of simple but distinctive visual clues to mediate the emotional and social state of a CVE user. Whilst currently concentrating on the face as the

channel for conveying emotions, the work has to be seen in a wider context in which the entire humanoid representation of a user can in principle act as the communication device in CVEs (see for example the work of Capin et al 1999 and Guye-Vuillème et al 1999 on gestures, Coulson 2002 on postures) and the experiments described here are therefore intended to set the foundation for further work on the expression of emotion and attitude through such a virtual embodiment.

2.3. EXPRESSION OF EMOTION IN THE FACE

As noted before, the human face plays a complex role in visual human communication. Producing (encoding) and recognising (decoding) distinct facial expressions constitute a signalling system between humans (Russell and Férnandez-Dols 1997). Surakka and Hietanen (1998) see facial expressions of emotion clearly dominating over vocal expressions of emotion, and Knapp (1978) generally considers facial expressions as the primary site for communication of emotional states.

Ekman et al (1972) found that there are six universal facial expressions, corresponding to the following emotions: *Surprise, Anger, Fear, Happiness, Disgust/Contempt,* and *Sadness.* The categorisation is widely accepted, and considerable research has shown that these basic emotions can be accurately communicated by facial expressions (Zebrowitz 1997, Ekman 1999). Indeed, it is held that production and, to an extent, recognition of these six emotions has an innate basis. They can be found in all cultures, and correspond to distinctive patterns of physiognomic arousal (Argyle 1994). These six facial expressions of emotion (together with the neutral expression) were employed in the experimental work in order to establish whether knowledge and expectations from the real world can be applied to the three-dimensional, modelled head in the virtual world.

Ekman and Friesen (1975, 1978) pioneered the analysis and categorisation of facial movements during the expressing of emotion, and constructed an atlas of the human face. The atlas depicts each of the six facial expressions of emotion and has formed the basis for numerous experiments in social psychology (Argyle 1994). The face is segmented into three areas: upper face (brows and forehead), mid-face (eyes, eyelids, and root of the nose), and lower face (mouth, nose, cheeks, and chin), these being the areas capable of independent movement (Ekman and Friesen 1975). The original atlas consists of a series of photographs of these three areas of the face, each photograph assigned to one of the six emotions. For each of the emotions, there is more than one photograph for at least one facial area. For example, in the emotion category *surprise,* there is only one distinctive brow/forehead segment and one distinctive eyes/eyelids segment, but four different possible atlas segments for the lower face.

To offer a comprehensive description of the visible muscle movement in the face, Ekman and Friesen (1978) established the Facial Action Coding System (FACS). This was informed by a major body of work and is based on highly detailed anatomical studies of human faces. A facial expression is a high level description of facial motions, which can be decomposed into certain muscular activities, i.e relaxation or contraction, called Action Units (AUs). FACS identifies 58 action units that separately, or in various combinations, are capable of characterising any human expression. An AU corresponds to an action produced by a single muscle or a group of related muscles. AU1, for example, is the *inner-brow-raiser* - a contraction of the *central frontalis* muscle, whereas AU7 is the *lid-tightener*, tightening the eyelids and thereby narrowing the eye opening.

FACS is usually coded from video or photographs, and a trained human FACS coder decomposes an observed expression into the specific AUs that occurred, their duration, onset, and offset time (Bartlett 1998). From this system, very specific details about facial movement for different emotional expressions of humans in the real world can be ascertained. For instance, the brow seems capable of the fewest positions and the lower face the most (Knapp 1978). Certain emotions also seem to manifest themselves in particular areas of the face. The best predictors for *anger* for example are the lower face and the brows/forehead area, whereas *sadness* is most revealed in the area around the eyes (Knapp 1978).

For the purpose of the intended experiments, FACS was adapted to generate the expression of emotions in the virtual face. The relevant action units were applied to the virtual head and their effectiveness in terms of cognition and acceptance were tested. Figure 1 shows some alternative expressions for the *anger* emotion category, together with the corresponding virtual head expressions used in the experiment. All photographs are taken from *Pictures of Facial Affect* (Ekman and Friesen 1972CD).

Figure 1: Pictures of Facial Affect showing variations of *Anger*, with corresponding virtual heads

Interest in modelling the human face has been strong in the computer graphics community since the 1980s. Platt and Badler (1981) developed the first muscle-based model of an animated face, using geometric deformation operators to control a large number of muscle units. Parke (1982) and Terzopoulos and Waters (1993) further developed this by modelling the anatomical nature of facial muscles and the elastic nature of human skin, resulting in a dynamic muscle model. These muscles were then mapped to FACS parameters to allow simple but very effective expression animation.

The approach chosen for these experiments was feature-based and therefore less complex than a realistic simulation of human physiology. It is argued that this is sufficient, and in fact preferable, as it allows the most distinctive and essential features of a facial expression to be established. The face model was based on the H-Anim (2002) specification developed by the international panel that oversees the Virtual Reality Modeling Language (VRML). H-Anim aims to define a humanoid as a basic set of segments and joints that can be displayed in any VRML browser and animated by any application.

The face model used was characterised by the following muscle groups and joints (from H-Anim): left eyeball, right eyeball, left eyebrow, right eyebrow, left upper eyelid, right upper eyelid, temporomandibular (for moving the jaw), and included an animation model for the human face. Clearly, these seven parameters do not allow representation of all possible facial expressions. However, it is not necessary for the entire set of FACS action units to be reproduced to achieve the level of detail envisaged for the current face model - there is evidence that the human perception system can recognise physiognomic clues, in particular facial expressions, from very few visual stimuli (Dittrich 1991), and the experimental head model was therefore designed to show merely these most distinctive facial clues.

3. EXPERIMENTAL STUDY

The purpose of the experimental work was to establish whether the reduced set of action units proposed is indeed sufficient to convey the six universal emotions on avatar faces in a collaborative virtual environment. The facial expressions of emotion were presented in two different ways: as natural photographs or as animated virtual heads. Within each of these two factors, there were seven sub-levels (the six universal expressions of emotion together with *neutral*). Each of the twenty-nine subjects in the repeated measures design was shown 28 natural photographs and 28 corresponding virtual head images, in a random order. Each of the six emotion categories was represented in 4 variations, together with 4 variations of the neutral

face. The variations were defined by differences in expression of the same emotion rather than by differences in intensity.

The subjects' performance and results during the experiments were logged by the same piece of software that presented the stimulus material. The data collected for each facial expression of emotion consisted of: type of stimulus material, expression depicted by each of the facial areas, emotion category expected, and emotion category picked by the subject

Figure 2: Recognition screen

The recognition screen (Figure 2) displayed the images and provided buttons to select an emotion category. A post-test questionnaire collected quantitative and qualitative data complementing the data collected during the recognition task.

Statistical analysis (Mann-Whitney) suggested that overall recognition rates for FACS photographs (78.6% overall) were significantly higher than those for virtual heads (62.2% overall).

However, this significant result was attributable totally to the *Disgust* category, which stood out as having a very low score for virtual faces (around 20%), while the result for photographs of disgust was over 70% (Figure 3). Excluding this category removed the significance of the overall difference. *Surprise*, *Fear*, *Happiness* and *Neutral* showed slightly better (but non-significant) results for FACS photographs, while *Anger* and *Sadness* virtual faces scored a little better (but not significantly so) than their FACS counterparts.

Thus, with the exception of the *Disgust* category, recognition was as successful with each virtual head as it was with the directly corresponding FACS photograph. The constructed avatar faces identified as being the most distinctive are shown in Figure 4.
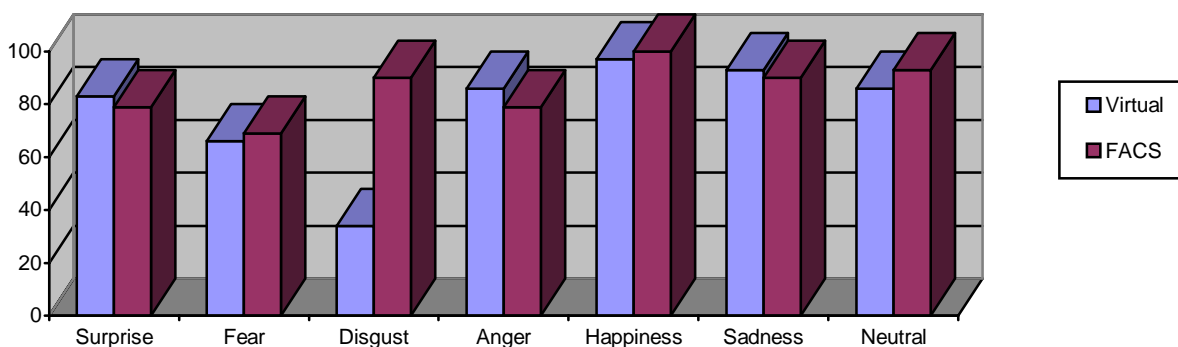

Figure 3: Summary of recognition rates

Recognition rates also varied significantly between subjects. The lowest scoring individual recognised 30 out of 56 emotions correctly (54%), the highest score was 48 (86%). Those who achieved better results did so homogenously between virtual and FACS images. Lower scoring subjects were more likely to fail at recognising virtual heads rather than FACS photographs.

5

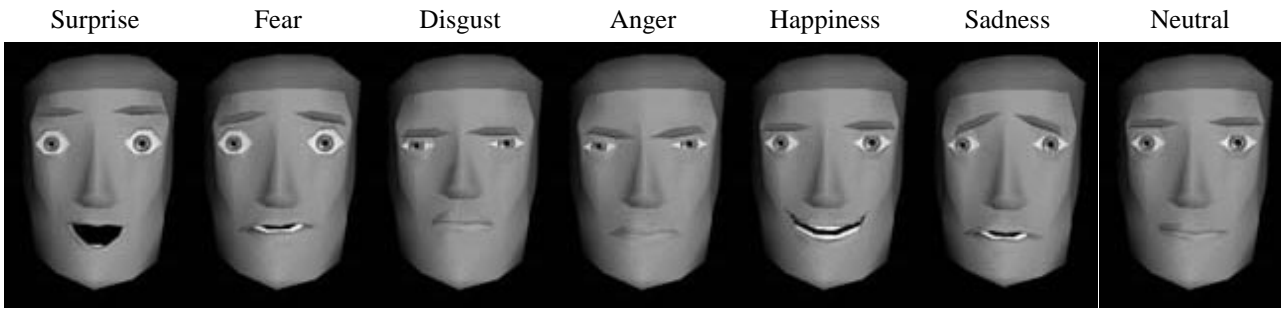| Surprise	Fear	Disgust	Anger	Happiness	Sadness	Neutral

Figure 4: Most distinctive facial expression in each category

The errors made by subjects when assigning expressions to categories are presented in (Figure 5). In particular, the matrix shows that the majority of confusion errors related to the category *Disgust*, an emotion frequently confused with *Anger*. When examining results for virtual heads only, anger (39%) was picked almost twice as often as disgust (22%). Further, with faces showing disgust, subjects often felt unable to select any given category and instead picked "Don't know", or suggested an alternative emotion. These alternatives were, for example, *aggressiveness*, *hatred*, *irritation*, or *self-righteousness*.

Emotion category	Responses [Virtual/FACS]							Other/Don't know
	Surprise	Fear	Disgust	Anger	Happiness	Sadness	Neutral	
Surprise	.67 .85	.06 .07	.00 .00	.00 .01	.23 .00	.00 .00	.01 .00	.03 .08
Fear	.15 .19	.41 .73	.00 .04	.30 .00	.03 .00	.03 .00	.02 .00	.06 .03
Disgust	.01 .02	.02 .00	.22 .77	.39 .14	.01 .00	.04 .00	.10 .01	.21 .07
Anger	.03 .04	.00 .04	.00 .03	.77 .72	.02 .00	.03 .03	.11 .05	.05 .09
Happiness	.01 .00	.01 .00	.01 .00	.01 .00	.64 .84	.03 .00	.26 .15	.04 .02
Sadness	.06 .00	.09 .10	.00 .00	.00 .01	.01 .01	.85 .66	.03 .09	.01 .07
Neutral	.03 .00	.03 .00	.01 .00	.00 .01	.00 .02	.11 .01	.78 .94	.04 .02

Figure 5: Error matrix for emotion categorisation *(Note: rows give per cent occurrence of each response. Confusion values above 10% are indicated yellow, above 20% orange, above 30% red)*

The error matrix further reveals that *Fear* was often mistaken for *Surprise*, a tendency that was also observed in several of Ekman's studies (1999). These two emotions share similar visual characteristics, and Ekman provides three sets of indicators to distinguish whether a person is afraid or surprised. However, all three involve context and timing of the fear-inspiring event – factors that are not perceivable from a still image. Poggi and Pelachaud (2000) and Bartneck (2001) similarly found that context and seeing the 'unfolding' of the emotion over time improved recognition rates. This suggests that in situations where the facial expression can be animated or displayed in context, recognition rates will be higher.

4. CONCLUSIONS AND FURTHER WORK

This study has shown that recognition is not guaranteed for all expressions, or all variations of a particular emotion category. Not surprisingly, the critical issues here are similar to those already identified by social psychology. Firstly, although accepted categories exist, emotions can vary in intensity and inevitably there is a subjective element to recognition. When modelling and animating facial features, such ambiguity in interpretation can be minimised by focussing on, and emphasising the most distinctive visual clues of a particular emotion.

Secondly, context plays a crucial role in emotion expression and recognition. Effective, accurate mediation of emotion is closely linked with the situation and other related communicative signals. A reliable interpretation of facial expressions cannot work independently of the context in which they are displayed. It

is anticipated that at least some confusion of emotions will be avoided when facial expression of emotion operates within the interactive animated setting of a CVE. Likewise, further work on emotion recognition in a real-time VR setting has to consider the effects timing has on display and interpretation of emotion. For example, showing surprise over a period of, say, a minute would, at the very least, send confusing or contradictory signals.

Finally, certain emotions were confused more often than others, most notably *Disgust* and *Anger*. This was particularly the case for virtual head expressions. Markham and Wang (1996) observed a similar link between these two emotions when showing photographs of faces to children. Younger children especially (aged 4-6) tended to group certain emotions together, whereas older children (aged 10+) usually seemed to differentiate correctly.

Nevertheless, this experimental work has provided strong evidence that creating virtual face representations based on the FACS model but using only a limited number of facial features does allow emotions to be effectively portrayed visually and gives rise to recognition rates that are comparable with those of the corresponding photographs. In consequence, the top-scoring expressions shown in Figure 4 may be taken to provide a sound basis for building emotionally expressive avatars to represent agents, as well as human users, in collaborative virtual environments.

5. ACKNOWLEDGEMENTS

6. REFERENCES

Argyle, M. (1994) *The Psychology of Interpersonal Behaviour (Fifth edition)*, London, Penguin Books

Argyle, M. (1988) *Bodily Communication (second edition)*, New York, Methuen & Co. Inc

Bartlett, M.S. (1998) *Face Image Analysis by Unsupervised Learning and Redundancy Reduction*, Ph.D. Thesis, University of California, San Diego

Bartneck, C. (2001) Affective Expressions of Machines, in *CHI 2001 Conference Proceedings*, Seattle, USA

Benford, S.D., Bowers, J., Fahlén, L.E., Greenhalgh, C.M., Snowdon, D. (1995) User Embodiment in Collaborative Virtual Environments, in *Proceedings of 1995 ACM Conference on Human Factors in Computing Systems (CHI'95)*, Denver/Colorado, ACM Press

Capin, T.K., Pandzic, I.S., Thalmann, N.M., Thalmann, D. (1999) Realistic Avatars and Autonomous Virtual Humans in VLNET Networked Virtual Environments, in *Virtual Worlds on the Internet*, Earnshaw, R.A., Vince, J. (eds.), IEEE Computer Science Press, ISBN 0818687002

Cooper, B., Brna, P., Martins, A. (2000) Effective Affective in Intelligent Systems – Building on Evidence of Empathy in Teaching and Learning, in Ana Paiva (Ed.) Affective Interactions: Towards a New Generation of Computer Interfaces, Lecture Notes in Artificial Intelligence 1814, ISBN 3-540-41520-3, Springer Verlag, pp 21-34

Coulson, M. (2002) Expressing emotion through body movement: A component process approach, in *Proceedings of AISB Symposium on Animated Expressive Characters for Social Interaction, London, UK, ISBN 1-902-95625-6, pp. 11-16*

Damásio, A.R. (1994) Descarte's Error: Emotion, Reason and the Human Brain, Avon Books, New York

Dittrich, W.H. (1993) Action categories and the perception of biological motion, in *Perception,* vol 22, pp. 15-22

Dittrich, W.H. (1991) Facial motion and the recognition of emotions, in *Psychologische Beiträge*, 33, 366-377

Dittrich, W.H., Troscianko, T., Lea, S.E.G., Morgan, D. (1996) Perception of emotion from dynamic point-light displays presented in dance, in *Perception*, vol 25, pp 727-738

Donath, J. (2001) Mediated Faces, in M. Beynon, C.L. Nehaniv, K. Dautenhahn (eds.), *Cognitive Technology: Instruments of Mind*, Proceedings of the 4th International Conference on Cognitive Technology, Warwick, UK, August 2001

Dumas, C., Saugis, G., Chaillou, C., Degrande, S., Viaud, M.L. (1998) A 3-D Interface for Cooperative Work, in *Proceedings of Collaborative Virtual Environments 1998 (CVE'98)*, June 1998, Manchester, UK

Durlach, N., Slater, M. (1998) Presence in Shared Virtual Environments and Virtual Togetherness, presented at *the BT Workshop on Presence in Shared Virtual Environments*, Ipswich, June 1998

Ellis, H. D. (1990) Developmental trends in face recognition, in *The Psychologist: Bulletin of the British Psychological Society* (3), pp. 114-119

Ekman, P. (1999) Facial Expressions, in T. Dalgleish and M. Power (eds.), *Handbook of Cognition and Emotion*, New York, John Wiley & Sons Ltd.

Ekman, P., Friesen, W.V., Ellsworth, P. (1972) *Emotion in the Human Face: Guidelines for Research and an Integration of Findings*, New York, Pergamon Press Inc.

Ekman, P., Friesen, W.V. (1975) *Unmasking the Face*, New Jersey, Prentice-Hall Inc.

Ekman P., Friesen W.V. (1975CD) *Pictures of Facial Affect*, CD-R, University of California, Department of Psychology, San Francisco, USA

Ekman, P., Friesen, W.V. (1978) *Facial Action Coding* System, Consulting Psychologists Press Inc.

Fleming, B, Dobbs, D. (1999) *Animating Facial Features and Expressions*, Charles River Media, Boston

Godenschweger, F., Strothotte, T., Wagener, H. (1997) Rendering Gestures as Line Drawings, in *Proceedings of International Gesture Workshop 1997 (GW97)*, Bielefeld, Germany, Springer Verlag, ISBN 3 540 64424 5

Guye-Vuillème, A., Capin, T.K., Pandzic, I.S., Thalmann, N.M., Thalmann, D. (1999) Nonverbal communication interface for collaborative virtual environments, in *Virtual Reality: Research, Development and Applications*, vol 4(1), pp. 49-59

Hindmarsh, J., Fraser, M., Heath, C., Benford, S. (2001) Virtually Missing the Point: Configuring CVEs for Object-Focused Interaction, in Churchill, Snowdon and Munro (eds.), *Collaborative Virtual Environments: Digital Places and Spaces for Interaction*, CSCW Series, ISBN 1-85233-244-1, Springer Verlag London, pp. 115-139

H-Anim (2002) *Specification for a Standard VRML Humanoid*, H-Anim Working Group, Online Document, URL http://www.h-anim.org

Kendon, A. (1983) Gesture and Speech: How they Interact, in *Nonverbal Interaction*, J.M Wiemann and R.P. Harrison (eds.), Beverly Hills, Sage Publications

Knapp, M.L. (1978) *Nonverbal Communication in Human Interaction (2*nd Edition) Holt, Rinehart and Winston Inc., New York

Lisetti, C.L., Douglas, M., LeRouge, C. (2001) Intelligent Affective Interfaces: A User-Modeling Approach for Telemedicine, in *Proceedings of International Conference on Universal Access in HCI (UAHCI)*, at HCI International, New Orleans, LA, August 2001, Elsevier Science Publishers B.V.

Lisetti C.L., Schiano, D.J. (2000) Facial Expression Recognition: Where Human-Computer Interaction, Artificial Intelligence and Cognitive Science Intersect, in *Pragmatics and Cognition* (Special Issue on Facial Information Processing), vol 8(1), pp. 185-235

Markham, R., Wang, L. (1996) Recognition of emotion by Chinese and Australian children, in *Journal of Cross-Cultural Psychology*, vol 27(5), pp. 616-643

McGrath, A., Prinz, W. (2001) All that Is Solid Melts Into Software, in Churchill, Snowdon and Munro (eds.), *Collaborative Virtual Environments - Digital Places and Spaces for Interaction*, CSCW Series, ISBN 1-85233-244-1, Springer Verlag London, pp. 99-114

Morris, D., Collett, P., Marsh, P., O'Shaughnessy, M. (1979) *Gestures, their Origin and Distribution*, London, Jonathan Cape Ltd.

Ødegård, O. (1993) Telecommunications and Social Interaction: social construction of virtual

Parke, F. (1982) *Parameterized modeling for facial animation*, IEEE Computer Graphics and Applications, vol 2(9), pp. 61-68

Platt, S.M., Badler, N.I. (1981) Animating facial expression, in *ACM SIGGRAPH Conference Proceedings*, vol 15(3), 245-252

Picard, R. (1997) Affective Computing, MIT Press

Poggi, I., Pelachaud, C. (2000) Emotional Meaning and Expression in Animated Faces, in Ana Paiva (Ed.) *Affective Interactions: Towards a New Generation of Computer Interfaces*, Lecture Notes in Artificial Intelligence 1814, ISBN 3-540-41520-3, Springer Verlag, 182-195

Russell, J.A., Férnandez-Dols, J.M. (1997) What does a facial expression mean?, in *The Psychology of Facial Expression,* J.A. Russell and J.M. Férnandez-Dols (eds.), Cambridge University Press

Stoney, S., Wild, M. (1998) Motivation and interface design: maximising learning opportunities, in *Journal of Computer Assisted Learning*, 14, 40-50, Blackwell Science Ltd.

Strongman, K.T. (1996) *The Psychology of Emotion. Theories of Emotion in Perspective (Fourth Edition)*, New York, Wiley & Sons

Surakka, V., Hietanen, J.K. (1998) Facial and emotional reactions to Duchénne and non-Duchénne smiles, in *International Journal of Psychophysiology*, 29, 23-33

Terzopoulos, D., Waters, K., (1993) Analysis and synthesis of facial image sequences using physical and anatomical models, in *Pattern Analysis and Machine Intelligence*, vol 15(6), pp. 569-579

Thalmann, D. (2001) The Role of Virtual Humans in Virtual Environment Technology and Interfaces, in R.A. Earnshaw, R.A. Guedj, J.A. Vince (Eds.) *Frontiers of Human-Centred Computing, Online Communities and Virtual Environments*, Springer Verlag London, pp. 27-39

Zebrowitz, L.A. (1997) *Reading Faces: Window to the Soul?*, Boulder, Colorado, Westview Press