

The Significance of Textures for Affective Interfaces

Paula M. Ellis and Joanna J. Bryson

Department of Computer Science, University of Bath, Bath BA2 7AY, United Kingdom
paulini13@hotmail.com, J.J.Bryson@bath.ac.uk

Abstract. This paper reports experiments demonstrating that the extent to which subjects ascribe emotions to VR faces is highly dependent on textures applied to the face. We demonstrate this for both a photo-realistic vs. non-photo-realistic texture pair and for a male vs. female texture pair. In both cases, experiments were conducted over the Internet on still frames taken from a well-controlled VR emotion modelling system. Given the enormous extent to which textures determine emotion recognition, we consider this a critical area for future research in affective virtual agents.

1 Introduction

There is much disagreement in the academic world regarding facial representation of emotion [6, 34]. There is also little available research regarding the success of virtual agents in effectively displaying their emotions to users. Haddad and Klobas [15] suggest that “character-agent visual representation” may influence the effectiveness of information delivery. Predinger and Ishizuka [23] argue that there is an ongoing debate of how to make agents more ‘life-like’, whether this is achieved by employing photo-realistic or more cartoon-like faces. As Predinger and Ishizuka suggest, it is the cartoon-like characters that tend to be more readily available in the entertainment/video game sector. They argue that users have higher expectations for the performance of the more realistic characters as opposed to the cartoon-like ones. Realistic characters take the risk that users may notice minor discrepancies in their actions that they might not notice in less realistic characters. Haddad and Klobas [15] go on to present evidence that in academic fields outside of the character-agent concern, the feeling also goes that less realistic is better. Graphic designers suggest that more cartoon-like animated characters more effectively convey information. Yet most VR literature points toward the supposed advantage of photo-realistic faces.

This paper reports our preliminary efforts to determine whether it is better to use photo-realistic or non-photo-realistic textures on VR faces designed to communicate emotion. While we do not yet have enough exemplar faces to determine that issue conclusively, we have shown highly significant results in the recognized emotion ascribed to identical models with different textures. This is true both for our photo-realistic vs. non-photo-realistic texture pair and for a male vs. female texture pair.

2 Background and Related Research

Our main aim is to discover on which type of facial skin texture an emotion can be most easily recognised. Past research appears inconclusive and contradictory in its findings.

Although Fabri et al. [9] found that participants were 78.6% successful at recognising the correct emotion on photographs but only 62.2% successful with the virtual heads, on closer inspection of their data it can be seen that fear and disgust had much lower scores for the virtual heads than for the other expressions and without these anomalies the results may have shown no significant difference. Much past research, however, has been concerned with how effective virtual faces have been in assisting the system they are attached to in conveying particular information. Very little research has concentrated on whether or not these faces can actually convey particular facial expressions to a recognisable degree.

2.1 What is Emotion?

Ferh and Russell [10, p. 177] observed that “everyone knows what emotion is until asked to give a definition” and Shaver et al. [25, p. 117] argue that “despite an enormous increase in research... there is still no widely accepted definition of emotion.” Izard [18], who reviewed the available literature on emotion in 1969, discovered that “the area of emotional experience... is one of the most confused and ill defined in psychology.” In fact Gaggioli et al. [14] have suggested that there are over 90 different definitions of emotion in existence in the scientific literature. Humans convey emotions to each other in many ways, through voice, body language and facial expression. As Cassell [3] explains “we make complex representational gestures with our hands, gaze away and towards each other... and use the pitch and melody of our voices” to communicate emotion. Cassell goes on to explain that non-verbal behaviours play an important role in the design of Embodied Conversational Agents, such as gesture, eye gaze and facial display.

Davis [6] suggests that there is a need to consider five ‘basic’ facial expressions of emotion. These are fear, anger, disgust, sadness, and happiness. Other authors too, such as Ekman et al. [8], have argued for the existence of basic emotions, but it seems that the main problem is that there is huge difficulty in establishing a way to truly reference each of the emotions that humans can recognise. It is true that as more research is carried out, experimenters are faced with the problem that although most people can recognise and identify various emotions, they are all so very subjective in their identification. This may be due to a number of factors, such as the ability of people to read body language and not just facial expressions. In summary there is huge disagreement in the academic field as to what emotion actually is, and how it is represented by the human body. This will pose problems for any researcher wishing to test or measure emotions in some way, as it is so hard to establish controls.

Collier [4, p.68] suggests that “one of the most persistent controversies among researchers... has revolved around the issue of whether facial expressions are learned or innate.” Fox [12] argues for three main perspectives that can help to explain the need for humans to recognise facial emotional expressions. He explains the biological, behavioural and cognitive arguments for the need to recognise expressions of emotion. The need to recognise the emotion another person is displaying is clearly an essential tool for life. It has been found that infants as young as 3.5 months old are able to recognise different emotions just by looking at still images of faces, supporting the argument that there is an evolutionary or behavioural advantage to emotion recognition [19]. Fox

et al. [11, p. 61] suggest that humans are “hard-wired for facial recognition, especially for the recognition of anger or threat.” Their experiment has shown that babies as young as five months old, can discriminate between the facial expressions of fear and anger. Another study by Hansen and Hansen [16] found that adult humans have the ability to spot an angry face in a crowd faster than a happy face [13, p. 94].

2.2 Emotional Expression in Virtual Agents

There has as yet been little work in the evaluation of animated virtual characters and their faces. One of the first virtual humans to be created was the ‘Boeing Man’. This was a three dimensional model, used to aid engineers in the building of cockpits. It was made using a collection of three-dimensional line segments, with articulated joints. There was only limited facial detail and no varying of the facial expression [21]. Massaro et al. [20] carried out an evaluation of a talking head named ‘Baldi’, which was constructed of polygons and approximately 900 surfaces joined together. Massaro et al. aimed to discover how informative certain properties of the Baldi’s was. The experiments concluded that participants were 94%, 95% and 73% correct at identifying happy, angry and sad faces, as demonstrated by Baldi.

Vinayagamoorthy et al. [32] found evidence that users in a virtual environment respond well to “humanoid representations of other users” in the environment. They aimed to test the importance of behavioural realism in avatars and virtual humans in virtual environments. Vinayagamoorthy et al. asked, should the behaviour of avatars mimic real life and to what extent? They also aimed to explore the optimum level of visual realism needed to make a character believable in a virtual environment. They reference the work of Strippgen [27], who has suggested that participants will expect more visually realistic avatars to behave “in a manner that portrays greater human like qualities.” In their study Vinayagamoorthy et al. [32] used a realistic avatar and a cartoonish avatar, which were modelled onto virtual faces that gave either realistic or non-realistic eye gaze. They found that the less realistic the avatar, the less effect the realism of eye movement had on the effectiveness of avatar communication. Others have supported these findings, suggesting that the employment of apparently human agents has only served to raise the users expectations of its performance, placing more pressure on the system to perform as ‘humanly’ as possible [24]. Takeuchi [29] suggested that when more realistic faces are used, the user will spend more time trying to interpret its expression that actually engaging in the task. De Rosis et al. [7] present a 3D embodied agent known as ‘Greta’. They argue that “the more a character aims at being realistic... the more complex its implementation becomes.”

Fabri et al. [9] developed a study that aimed to show whether photo-realistic or animated facial expressions could be understood more easily. The photo-realistic expressions were represented using photographs and the more animated cartoon facial expressions were displayed using virtual heads. Seven different facial expressions were represented by each of the two facial forms, in four variations of expression. Software was used to present each of the 56 facial images to participants in a random order and this software also recorded the emotion that each face was assigned by the participants. Mann-Whitney statistical analysis was carried out on the data and the findings

suggest that realistic photographs are easier to interpret than the expressions of the virtual heads. Participants were 78.6% successful at recognising the correct emotion on the photographs but only 62.2% successful with the virtual heads. However, on closer inspection of the data it can be seen that fear and disgust had much lower scores for the virtual heads than for the other expressions and without these anomalies it may be possible that results would have shown no significant difference

Ward et al. [33] carried out an investigation monitoring the facial movements of participants engaged in a web-based task. Participants were asked to complete an online test that contains one surprise event and FaceStation 1.2, facial tracking software was used to track their faces as they completed the task. The captured data from this was then transferred to create the movements of a virtual face. ‘Judges’ were asked to observe the footage of the real faces completing the tasks and of the virtual faces that had been created by the tracking software [33]. With respect to the reactions of both types of face judges considered the movements of the virtual and real faces to be in agreement for 7 out of the 15 pairs of faces. In their conclusion Ward et al. 2003, suggest that the failure of the facial tracking software is to blame for the results. However, in presenting their work, Ward et al. suggested that participants in some cases seemed to be recognizing emotions in the virtual faces constructed by the facial tracking system that they had failed to notice on the films of real human faces [2]. This report, in addition to some similar experiences with demonstrating the DER [31, see below] in different texture conditions, lead to the present research.

3 Experiment 1

Our experiments use the Dynamic Emotional Representation (DER) facial representation model, developed by Tanguy et al. [31], to create faces representing emotions. Pairs of skin textures are tested showing various degrees of various emotions in experiments conducted over the Internet. Note that the values for varying degrees of ‘intensity’ are a somewhat arbitrary amount that is only relevant to the piece of software used for these experiments. However, they were entirely consistent across textures, which is what matters for the significance of our results.

3.1 Equipment and Stimuli

Tanguy et al. [31], present a facial animation platform with an integrated Dynamic Emotional Representation. The design of the DER is based on the Sloman [26] model of emotion representation, describing three-layer architecture to emotion. The face is animated using the Parke and Waters [22] abstract model for the facial mesh [30]. Tanguy et al. [31] explain that the DER is designed to “enforce consistency in the production of emotional facial displays,” providing, a “rich, real time representation of emotions. . . without their automatic generation.” The DER interface can be used to produce various facial expressions on the various skin textures that are available, such as a photo-realistic male face and a more cartoon-like male face. These skin textures are used to generate the facial images for the experiments. The DER package contains a user interface that allows users to select various facial skins and to view in real time

the changes of the facial components as an emotion is selected. The required facial skin texture is displayed to the user by using the interface to open the required file. Once a facial skin texture has been loaded, the user can select the 'Present of Textures' button that removes the photo-realistic skin texture from the face and displays a more simple animated or cartoon-like face. The textures used in our study can be seen in Fig. 1. Some of the related controls can be seen in Fig. 2.

The DER has been set to produce facial representations of the following expressions: 'Happy', 'Sad', 'Angry', 'Surprised' and 'Disgusted'. Past studies have shown that people often have trouble correctly distinguishing between various expressions. Hara and Kobayashi [17] found participants identifying emotions on a robot face would often confuse expressions of fear with surprise and disgust. Fabri et al. [9] also found that disgust was generally harder to identify on virtual heads. It is with this in mind that expressions of disgust are disregarded for this particular experiment.

For this experiment 24 still facial images were created. 12 still images are required for both facial types; the photo-realistic one (known as 'PR') and the cartoon-like one (known as 'An'). The user interface to the DER allows the user to select the desired emotion by clicking on a button labelled with the desired expression. When the button is clicked the skin displays the emotion and then returns to a neutral expression. It is important to check that various other settings are standardised for each face before aiming to capture a still image. As the DER, by its very nature, creates dynamic facial images, it is important that all other expression intensity values except the one being created are set to '0', or this will affect the way an expression is produced. See Figure 1.




Fig. 1. Virtual faces as created by the DER. (a) Photo-realistic male face, used in both experiments. (b) Non-photo-realistic face, used in Experiment 1. (c) Photo-realistic female face, used in Experiment 2.

It should be noted at this point that *intensity* in this context is an arbitrary amount because there is no way to measure the intensity of a facial expression. It is still useful however because the intensity measure serves as a way to express the difference between various emotions and to check that others are displayed to the same degree as another. The purpose of other settings is irrelevant to this study, but as each facial still is captured, it is critical to ensure that the settings for ‘release’ and ‘contraction’ are set to the same amount for every face. It is also essential to ensure that the various ‘tick boxes’ are in the same status for every face. See Figure 2.

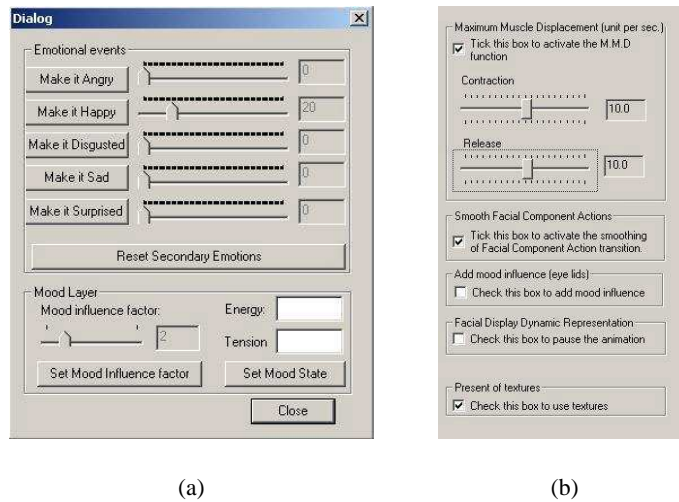


Fig. 2. (a) Shows the box that allows the user to select emotions and demonstrates how the intensity for all expressions other than happy are set to 0. (b) Shows an example of how the tick boxes and settings can be set up.

As mentioned above, 24 faces are produced. Both facial skin textures are captured displaying the four expressions: Happy, Sad, Angry and Surprised. These expressions are displayed in the intensities 10%, 30% and 50%. As each face is produced, screen shots are taken and then imported into a suitable package such as Microsoft Paint. The facial image can then be saved in jpg format ready to be called by the PHP code.

3.2 Design and Method

The experiment uses a repeated measures, forced response design, where all participants are exposed to the Dependent Variable being the various representations of facial expression. The forced responses participants must choose from are: ‘happy’, ‘sad’, ‘angry’, ‘surprised’ or ‘unsure’. A similar methodology was used by Breazeal [1], who

used a forced-choice design during evaluation of Kismet, a robot head. There participants were given the choice of ten labels and were asked to assign these to a number of still images of Kismet performing various facial expressions. Calder (2001) also used a forced choice design when evaluating animated humanoid heads. The repeated measures design also eliminates the possibility that any differences in results are caused by differences between participants, as all participants are exposed to all changes in the dependent variable [5].

In our study, faces are displayed to participants in a random order as generated by the PHP script, and no one face can be displayed more than once to each participant. Participants are directed to the experiment home page. Here they are able to read the experiment briefing and instructions. Participants are not asked for any personal information, but are asked to indicate the most truthful answer as they complete the experiment. When participants click on the start button they are directed to the first face, and prompted to select the button reflecting the emotion that they feel the face is showing. Answers are stored in the correct table for later analysis.

Participants were collected via an on-line chat forum frequented by the experimenter. A message on the forum asked participants to partake in the experiment, and informed them that no personal data would be collected. The only ethical considerations for this experiment are to inform participants that no personal data will be collected from them and ensure that they are provided with full contact details of the experimenter.

3.3 Results

Correct Expression	Texture Type				Significance
	Photo-Realistic correct of 90	percent correct	Non-Photo-Realistic correct of 90	percent correct	
Happy	86	95.6%	85	94.4%	N/S
Sad	59	65.6%	31	34.4%	$***\chi^2(1, N = 90) = 17.42, p < .001$
Angry	53	58.9%	24	26.7%	$***\chi^2(1, N = 90) = 19.09, p = .001$
Surprised	42	46.7%	40	44.4%	N/S
Totals	240	66.7%	180	50%	$***\chi^2(1, N = 90) = 20.57, p < .001$

Table 1. Total number and percentage of correct assignments for each group of expression.

Table 1 shows the results for the first experiment. We can see that, at least for these two exemplars, it is significantly easier ($\chi^2(1, N = 90) = 20.57, p < .001$) to correctly identify the emotion for the face with the photo-realistic texture than for the animated-style one. Note that there is no significant difference for the happy or surprised conditions, but all the variance results from the sad and angry conditions, both of which favour the photo-realistic face. No general conclusions about recognizing specific emotions can necessarily be drawn from this data, given the fact that the emotion

exemplars and what it means to be at 10, 20 or 30% of them have been set by hand by Tanguy and have not yet been thoroughly tested. However, to first approximation they do seem to be reasonably good models and well matched.

Regardless of the validity of the between-emotion measurements, the between-face measurements clearly show a significant, texture-dependent effect.

4 Experiment 2

As a first step to creating additional stimuli, we used one of the author’s face to create a female photo-realistic face. Because the framework underlying the face is still identical to the previous experiment, the result was somewhat androgynous. Out of interest, we decided to conduct a second experiment comparing the male and female textures.

4.1 Equipment, Stimuli, Design and Method

The equipment and stimuli were all largely as before, except for the different texture, which can be seen in Figure 1 as well.

The apparatus in this case consists of a simple piece of HTML and PHP code similar to that used in the main experiments and four sets of male and female faces showing the same intensity of the same expression. This code creates four basic web pages that display one pairing of a male and female face showing the identical expression, at a time. Participants are asked to select a single radio button, to log which face they believe is showing the stronger emotion. The page title says “Experiment on Emotions in Faces.” The text under the two (happy) stimuli reads “Please Select an option from the list below to indicate which face you believe looks more Happy. If you think there is no difference please select the ‘No difference’ button.” The options are ‘Male’, ‘Female’ or ‘No difference’. The participants for this experiment were again collected via an online forum, though less effort was made to recruit since, at the time, this experiment seemed less central to our study.

4.2 Results

Expression	Total number of each answer assigned by participants			Significance
	Male	Female	No Difference	
Happy	4	9	10	N/ S
Sad	1	21	1	$***\chi^2(2, N = 23) = 42.17, p < .001$
Angry	18	2	3	$***\chi^2(2, N = 23) = 31.43, p < .001$
Surprised	3	12	8	$*\chi^2(2, N = 23) = 7.9, p < .05$
Total	26	44	22	$**\chi^2(2, N = 23) = 13.43, p < .01$

Table 2. Participants’ assesment of the relative intensity of emotional expression with male vs. female textures. Note that, in fact, the underlying facial structure and intensity of expression were identical in all cases, so the expected choice should be no difference.

In experiment 2, we again see that assessed expression difference for two different textures is highly significant, although in this case which face is perceived as more emotional depends on the particular emotion expressed. Although there is a weakly significant trend toward ascribing more emotional expression to the female face, what we see is a strong gender effect in line with previous psychological results: participants are more sensitive to anger in the male face but to sadness and surprise in the female face. The strength of the significance of these results is particularly stunning given the fact that the female texture is not overtly female (e.g. is not wearing significant makeup) and, as mentioned earlier, is placed on a ‘male’ facial mesh identical to that of the male stimuli. However, it was labelled as ‘female’ in the answer key, which may in itself introduce some bias [24]. In future work we would like to re-run this experiment labelling the faces simply ‘a’ and ‘b’.

5 Discussion and Future Work

As we said from the outset, these results can only be seen as preliminary pilot work in so far as they represent the issues of photorealism vs. caricature or male vs. female stimuli. In particular with the photorealism issue, our non-photo-realistic texture is not particularly caricature-like. It is basically androgynous and lacks the highly caricatured, exaggerated features, such as big bushy eyebrows, a large nose or an over-exaggerated smile one expects in cartoon-like faces. The default texture for the DER’s frame is mainly made of pink pixels, with a white line that represents teeth and blue eyes. Although there are definite areas that represent bone structure above the eyes and around the cheeks, in general when looking at this facial representation, one is looking at a large amount of pink pixels. As these pixels move over a virtual bone structure to simulate the wrinkles formed by a smile, for example, the observer has trouble noticing such a movement as pixels of one colour are just moving around the facial texture. If more obvious features were added in then it is possible the facial texture would be much easier to interpret.

At the time these experiments were first conducted, we suggested also that the eyes were not sufficiently expressive. Collier [4] suggested that when conversing, humans spend most of their time looking at one another’s eyes. Sullivan and Kirkpatrick [28] also observed the areas of the face that we watch when conversing with each other. They found that for expressions of happiness, sadness and surprise, more time is spent looking at the mouth but the eyes are important when observing expressions of anger. The point to be made here is that clearly various parts components of the face are important when reading facial expressions. It is possible that the way the DER and facial textures work to portray facial emotions is not close enough to that of a real human face. This is in fact an area that has since been developed by Tanguy, and is still under development.

As for our own work, it is obviously essential to begin working with more mesh frameworks and a vastly greater number of texture stimuli in order to explore both the questions we have opened up. It would also be more useful to log more information about the participants, such as sex, age, cultural background and even ethnicity, to see whether these characteristics play a role.

Nevertheless, even these preliminary results give a clear warning to those working on developing virtual affective interfaces. The extent to which their emotional interface will be perceived as such is heavily determined by the texture they choose to apply to their characters.

6 Conclusions

We have shown that textures overlying the framework of a VR face have an enormous impact on how emotions are perceived on that face. In a series of experiments all using the same emotional expressions and the same mesh framework underlying the texture, we have shown highly significant differences in whether the intended emotion is perceived (on photo-realistic vs. non-photo-realistic textures) and even on the extent to which an expression is being displayed on side-by-side comparisons (on male vs. female textures). We recommend a great deal of further research is necessary in this area, as such significant differences from texture could have an enormous impact on the utility of any affective display.

Acknowledgements

Thanks to Emmanuel Tanguy and Robin Hodges for their assistance in running this project, and to our very helpful reviewers. This work was funded in part by The Engineering and Physical Sciences Research Council (EPSRC) Grant GR/S79299/01.

References

- [1] Cynthia Breazeal. Emotion and sociable humanoid robots. *International Journal of Human Computer Interaction*, 59:119–155, 2003.
- [2] Joanna Bryson. Personal communication, based on data presented at *HCI 2003* by Ward, Bell & Marsden, 2004.
- [3] Justine Cassell. Nudge nudge wink wink: Elements of face-to-face conversation for embodied conversational agents. In Justine Cassell, Joseph Sullivan, Scott Prevost, and Elizabeth Churchill, editors, *Embodied Conversational Agents*, pages 1–27. MIT Press, London, 2000.
- [4] Gary Collier. *Emotional Expression*. Lawrence Erlbaum Associates, London, 1985.
- [5] Hugh Coolican. *Research Methods and Statistics in Psychology*. Hodder & Stoughton, London, second edition, 1994.
- [6] Darryl N. Davis. Agents, emergence, emotion and representation. In *IECON '00: 26th Annual Conference of the IEEE Industrial Electronics Society*, volume 4, pages 2577–2582. IEEE, October 2000.
- [7] Fiorella De Rosis, Catherine Pelachaud, Isabella Poggi, Valeria Carofiglio, and Berardina De Carolis. From greta's mind to her face: Modelling the dynamics of affective states in a conversational agent. *International Journal of Human Computer Studies*, 59:81–118, 2003.

- [8] Paul Ekman, Wallace Friesen, and Phoebe Ellsworth. *Emotion in the Human Face: Guidelines for Research and an Integration of Findings*. Pergamon Press, New York, 1972.
- [9] M. Fabri, D. J. Moore, and D. J. Hobbs. Expressive agents: Non-verbal communication in collaborative virtual environments. In *Proceedings of AAMAS 2002 Workshop: Embodied Conversational Agents: Let's Specify And Evaluate Them!*, pages 402–409, July 2002.
- [10] B. Ferh and J. A. Russell. Concept of emotion viewed from a prototype perspective. *Journal of Experimental Psychology: General*, 113:464–486, 1984.
- [11] Elaine Fox, Victoria Lester, Riccardo Russo, R. J. Bowles, Alessio Pichler, and Kevin Dutton. Facial expression of emotion: Are angry faces detected more efficiently? *Cognition and Emotion*, 14(1):61–92, 2000.
- [12] Jeremy Fox. Factors of emotion recognition in faces: Three perspectives. *Journal of Young Investigators*, (3), March 2004.
- [13] Alan J. Fridlund. The behaviour, ecology and sociality of human faces. In Margaret S. Clark, editor, *Emotion*, pages 175–212. Sage, London, 1992.
- [14] A. Gaggioli, F. Mantovani, G. Castelnuovo, B. Wiederhold, and G. Riva. Avatars in clinical psychology: A framework for the clinical use of virtual humans. *CyberPsychology and Behavior*, 6(2):117–125, 2003.
- [15] Hanadi Haddad and Jane Klobas. Expressive agents: Non-verbal communication in collaborative virtual environments. In *Proceedings of AAMAS 2002 Workshop: Embodied Conversational Agents: Let's Specify And Evaluate Them!*, July 2002.
- [16] C. H. Hansen and R. D. Hansen. Finding the face in the crowd: An anger superiority effect. *Journal of Personality and Social Psychology*, 54:917–924, 1988.
- [17] Fumio Hara and Hiroshi Kobayashi. A face robot able to recognise and produce facial expression. In *IEEE International Conference on Intelligent Robots and Systems (IROS'96)*, Osaka, 1996.
- [18] C. E. Izard. The emotions and emotion constructs in personality and culture research. In R. B. Cattell, editor, *Handbook of Modern Personality Theory*. Aldine Press, Chicago, 1969.
- [19] R Kahana-Kalman and A. S. Walker-Andrews. The role of person familiarity in young infants perception of emotional expressions. *Child Development*, 72:352–369, 2001.
- [20] Dominic. W. Massaro, Michael. M. Cohen, Jonas Beskow, and Cole Ronald. A. Developing and evaluating conversational agents. In Justine Cassell, Joseph Sullivan, Scott Prevost, and Elizabeth Churchill, editors, *Embodied Conversational Agents*, pages 287–317. MIT Press, London, 2000.
- [21] Frederic I. Parke. Computer graphic models for the human face. In *Proceedings of the IEEE Computer Society's Third International Computer Software and Applications Conference (COMPSAC'79)*, pages 724–727, Chicago, November 1979.
- [22] Frederic I. Parke and Keith Waters. *Computer Facial Animation*. A. K. Peters Ltd., 1996.
- [23] Helmut Predinger and Mitsuru Ishizuka, editors. *Life-Like Characters: Tools, Affective Functions and Applications*. Springer, Berlin, 1998.
- [24] B. Reeves and C. Naas. *The Media Equation: How People Treat Computers, Television and New Media Like Real People and Press*. Cambridge University Press, 1996.

- [25] Phillip R. Shaver, Shelley Wu, and Judith C. Schwartz. Cross cultural similarities and differences in emotion and it's representation: A prototype approach. In Margaret S. Clark, editor, *Emotion*, pages 175–212. Sage, London, 1992.
- [26] Aaron Sloman. Beyond shallow models of emotions. *Cognitive Processing*, 2(1): 177–198, 2001.
- [27] Simone Strippgen. INSIGHT: A virtual laboratory for design, test and evaluation of autonomous agents. In Richard N. Zobel and Dietmar P. F. Möller, editors, *Twelfth European Simulation Multiconference: Simulation — Past, Present and Future*, Manchester, UK, June 1998.
- [28] L. A. Sullivan and S. W. Kirkpatrick. Facial interpretation and component consistency. *Genetic, Social and General Psychology Monographs*, 122(4):389–404, 1996.
- [29] A. Takeuchi. Situated facial displays: Towards social interaction. In *Proceedings of CHI '94*, pages 450–454. ACM Press, 1994.
- [30] Emmanuel Tanguy. An abstract muscle model for facial animations. Dissertation about a year project carried out during the third year of a computer science degree in the Univeristy of Sheffield., May 2001.
- [31] Emmanuel A. R. Tanguy, Philip J. Willis, and Joanna J. Bryson. A layered Dynamic Emotion Representation for the creation of complex facial animation. In Thomas Rist, Ruth Aylett, Daniel Ballin, and Jeff Rickel, editors, *Intelligent Virtual Agents*, pages 101–105. Springer, September 2003.
- [32] V. Vinayagamoorthy, M. Garau, Steed, and M. Slater. An eye gaze model for dyadic interaction in an immerse virtual environment: Practice and experience. *Computer Graphics Forum*, 23(1):1–11, 2004.
- [33] Robert Ward, Dennise Bell, and Phil Marsden. A exploration of facial expression tracking in affective HCI. In Eamonn O'Neill, Philippe Palanque, and Peter Johnson, editors, *People and Computers XVII — Designing for Society (HCI '03)*, pages 383–399, Bath, September 2003. Springer.
- [34] Michelle S. M. Yik and James A. Russell. Interpretation of faces: Cross-cultural study of a prediction from Fridlund's theory. *Cognition and Emotion*, 13(1):93–104, 1999.