Using a Human Face in an Interface

Janet H. Walker[1]
Cambridge Research Lab
Digital Equipment Corporation
One Kendall Square, Bldg 700
Cambridge, MA 02140 USA
E-mail: jwalker@world.std.com

Lee Sproull
Department of Information Systems
School of Management
Boston University
704 Commonwealth Avenue Suite 400
Boston, MA USA

R. Subramani
Department of Information Systems
School of Management
Boston University
704 Commonwealth Avenue, Suite 400
Boston, MA USA

ABSTRACT

We investigated subjects' responses to a synthesized talking face displayed on a computer screen in the context of a questionnaire study. Compared to subjects who answered questions presented via text display on a screen, subjects who answered the same questions spoken by a talking face spent more time, made fewer mistakes, and wrote more comments. When we compared responses to two different talking faces, subjects who answered questions spoken by a stern face, compared to subjects who answered questions spoken by a neutral face, spent more time, made fewer mistakes, and wrote more comments. They also liked the experience and the face less. We interpret this study in the light of desires to anthropomorphize computer interfaces and suggest that incautiously adding human characteristics, like face, voice, and facial expressions, could make the experience for users worse rather than better.

KEYWORDS

User interface design, multimodal interfaces, anthropomorphism, facial expression, facial animation, personable computers

INTRODUCTION

Humanizing computer interfaces has long been a major goal of both computer users and HCI practice. Humanizing has at least two aspects, that of making interfaces easier and more comfortable to use (e.g., [22, 28]) and of making interfaces more human-like. Anthropomorphizing the interface entails adding such human-like characteristics as speech output (e.g., [11]), speech recognition (e.g., [16]), auditory and kinesthetic feedback (e.g. [14, 30]), models of human discourse [26, 29] and emotion [24], and social intelligence [5, 25].

CHI94-4/94 Boston, Massachusetts USA
© 1994 ACM 0-89791-650-6/94/0085...$3.50

Previous research indicates that simply adding human qualities to an interface is not in itself guaranteed to produce a better human-computer interaction experience. For example, "human-like" error messages, in contrast with "computer-like" error messages, produced more negative criticisms from users [25]. In a study in which subjects were tutored and evaluated by computer program, a voice interface produced more negative criticisms than a text interface [20]. In a historical hypertext database using icons of historical characters as guides, users over generalized from the character icons, expecting them to show personality, motivation, and emotion [23].

The most human-like interface of all of course would embody human characteristics in an agent with human form, of which the human face is one of the most compelling components. Infants are born with information about the structure of faces; at birth infants exhibit preference for face-like patterns over others [6]. By the age of two months infants begin to differentiate specific visual features of the face [19] and process facial expressions [8]. Faces can induce appropriate behavior in social situations; covering faces with masks can produce inappropriate behavior [9]. Faces, particularly attractive ones, even sell soap; physically attractive models are found to be effective in improving peoples' responses to advertisements [4].

There is some history of using face icons and faces in interfaces [18, 31, 32]. Perhaps the most famous example of an embodied agent is "Phil," the computer-based agent played by a human actor, which appeared in the Knowledge Navigator envisionment videotapes from Apple Computer [1].

Exploration of human-like interfaces has been limited to date by technology but this situation is changing rapidly. The base technology needed to implement a variety of personable interfaces is at, or close to, commercial availability. With a combination of speech synthesis technology (commercially available) and facial animation (in research prototype), it is possible to display a synthetic talking face on a workstation screen [33, 34]. The face is

1 Author's current address: 103 Raymond Street, Cambridge MA 02140 USA; email: jwalker@world.std.com.

an image of a human face with the mouth animated in synchrony with speech output from a text-to-speech algorithm. The face can speak arbitrary text and could participate more or less fully in an interaction with a user (depending of course on the underlying programming). At one extreme, the face could simply provide a stylized greeting and introduce the user to a more conventional interface. Moving toward the other extreme, the face could represent the computer side of an entire interaction, speaking all words that would otherwise be displayed on the screen as text and responding to the user orally instead of via text.

Talking faces may be particularly problematic in interfaces precisely because the human face is such a powerful stimulus. A talking face could focus people's attention on the face itself rather than on the task they are trying to perform, thereby leading to lower performance. Alternatively, it might cue reminders of social evaluation, thereby leading to higher performance. A talking face could also engender strong expectations for social intelligence from the interface, thereby leading to confusion or frustration if those expectations are not met. This paper reports the results of an initial investigation of reactions to a talking face.

RESEARCH QUESTIONS

The first question is, simply, are people willing to interact with a talking face? It is possible that people would find the prospect so bizarre that they would refuse to participate in a computer interaction in which the computer side of the interaction was represented by a talking face.

The second question is, even if they are willing to participate, will they be so distracted that their performance is seriously degraded?

Finally, the most important question is how people experience the interaction with the face. How human does it seem? Does it evoke a social response from the user?

We investigated these questions in an exploratory study using the social context of the interview survey. People are quite familiar with the general social structure and form of interview surveys. One person asks questions, usually through an agent such as a questionnaire or interviewer; and another person answers them. People are accustomed to answering questions asked by such agents and there is an extensive literature on how the nature of the agent affects peoples' responses. See, for example, Bailey, Moore, and Bailar [3] for the effects of interviewer demeanor on the obtained data, Schuman and Presser [27] for the effects of question wording and order on the obtained data. Generally, surveys delivered by human agents (face-to-face or telephone interviewers) are more socially involving than those delivered by paper and pencil. Thus response rates are higher and missing data rates are lower. But social involvement can also lead to social posturing; hence surveys delivered by human agents elicit more biased reports of socially desirable and undesirable behavior (e.g., [7]).

For exploring the first two issues, we compared text delivery of survey items to having them spoken by the synthetic face. To look specifically for social response effects and to control for simple differences due to the method of delivery, we compared two versions of the speaking face that differed in expression.

METHODS

The experimental context was a computer-based survey in which subjects received questions in either text or spoken form and typed in their answers. We used a between-subjects design with subjects assigned randomly to one of the three presentation conditions: questions spoken by a face with a neutral expression, spoken by a face with a stern expression, or text only.

Subjects

The subject population was the staff of a computer research laboratory in a large industrial corporation. This included full-time research, support, and administrative staff, part-time research staff, external consultants, and some off-site people associated with the laboratory on a permanent basis (for example, financial analyst, personnel consultant). People who worked in the computer support organization or who were involved in conducting the study were excluded. The population thus defined consisted of 49 people, who were assigned randomly to the three conditions. We checked to be sure that part-time staff were not over-represented in any one condition.

Task

All subjects answered questions designed to measure user satisfaction with the computer support services in the laboratory. The survey had been commissioned by the computer support organization. Thus the content of the questionnaire was both realistic and salient to the respondents. Subjects were informed of the purpose of the study and were assured that their identities would be concealed in the final report.

Questionnaire

The questionnaire contained 79 items: four background questions, 59 questions about computing attitudes and behaviors, and 16 questions about the experience of participating in the study. The computing attitudes questions were a mix of open-ended questions (for example, what hardware do you have in your office?) and fixed-response scale questions (for example, on a scale from 1 to 10, how easy is it for you to talk with the support staff about your needs?). The fixed-response questions were based on a modified version of Zeithaml *et al*'s service quality questionnaire [36, 37].

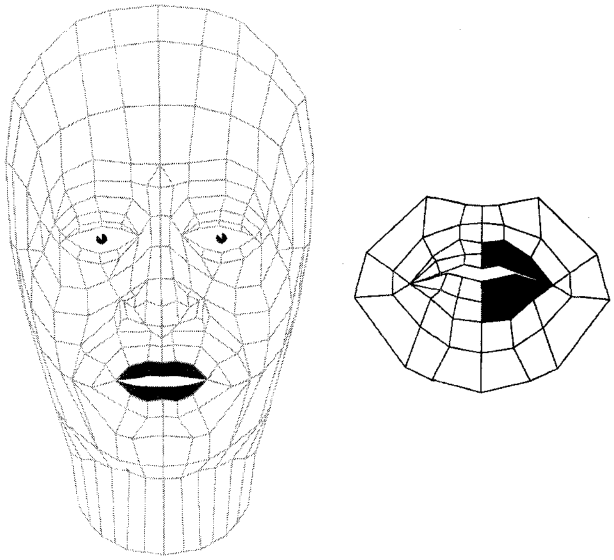

Figure 1. (a) Geometric model underlying the animation (b) Viseme (mouth posture) corresponding to the phoneme *r* in the word *red*.

The questions about the experience of answering questions consisted of eleven 10-point scale questions assessing "the question asker" (for example, how happy does the question asker seem?), four 10-point scale questions assessing subjects' reactions to the experience itself (for example, how comfortable were you in answering these questions?), a time estimation question, and an open-ended question inviting comments about the experience.

Procedure

Each subject completed the study individually in one of the offices in the laboratory, equipped with a computer workstation. The subjects used workstations regularly as part of their jobs and the system used in the study presented an environment familiar to most of them. The experimenter introduced the survey and explained how the questions would be displayed. Each subject had the opportunity to answer three practice questions to make sure they understood how to control the interface before the experimenter left the room. The survey was self-paced; subjects were free to work on it as long as they wished.

Subjects in the text condition first read an introduction to and explanation of the computer support survey in a text window. They then saw the questions displayed, one at a time, in that window. They typed their answer after each question. Subjects in the face conditions heard rather than saw the same introduction and questions, and typed in their responses in the same way as did subjects in the text condition. The face remained on the screen between questions. At any point, subjects could scroll backward to see any prior question or edit any answer. For all conditions, a help window about controlling the interface remained on the screen at all times.

Apparatus

The workstation in the office was a Digital Equipment Corporation DECstation 5000/200, equipped with a 21-inch color monitor, a DECaudio board, amplifier, and speakers. Images were shown in grey scale, not in color. All materials were pre-computed to achieve acceptable performance and stored on a local disk to prevent variability in display speed due to network traffic. The experimental sessions were managed using the Lisp facilities of Gnu Emacs.

The face was produced by texture-mapping an image onto a geometric wire-frame model (shown in Figure 1a; see [33] for details). On the screen, the face occupied 512 x 320 pixels, which is about half life size. This was the maximum size provided by the technology. The mouth was animated by computing the current mouth posture (viseme, shown in Figure 1b) corresponding to the current linguistic unit (phoneme) and applying cosine-based interpolation between the images [34]. The voice was produced by a software implementation of the KLSYN88 revisions of the DECtalk text-to-speech algorithm [17]. The DECtalk parameters used were the "B" voice, a fairly neutral voice in the female pitch range, at 160 words a minute. DECtalk speech is acceptably comprehensible at this rate [10]. The synchronized face and voice were presented using the Software Motion Pictures component of DECmedia [21].

The neutral and stern expressions appear in Figure 2. The

Figure 2. Texture mapped face with (a) neutral expression and (b) stern expression.

neutral expression (Figure 2a) was selected from a videotape of a female speaker and converted to the form required for texture mapping. The stern expression (Figure 2b) was synthesized from the neutral expression by contracting the corrugator muscles in the underlying physical model for the animation, thus pulling the inner portion of the eyebrows in and down. The expression produced by contracting these muscles is recognized as conveying negative emotion such as anger and threat [2, 12]. The mouth itself was not involved in forming the expression; subjects in both face conditions saw exactly the same lip animation and heard the same voice.

After completing the study, seven people from the subject population rated both of the face images on a 30-item set of bi-polar adjectives used to make judgments relevant to emotion and personality (taken from [13, 15]). The face we called "stern" was rated consistently more negatively by these subjects, thereby validating our label for the face.

Dependent Measures

Attitude measures toward computing resources and staff were defined as the sum of the responses to the ten 10-point scale items on these topics. Overall measures of response extremity were defined as the frequency of using the extremes on the 10-point scales in the body of the questionnaire (37 items). Responses to the open-ended question asking "how this method of asking questions could be improved" were coded independently by two coders. Responses were first divided into remarks and then coded for comments about the face, voice, text editor, and question wording. Comments were also coded for general positive or negative affect and for the presence of first person pronouns (as an indicator of involvement). The system logged how long each session lasted and how many words subjects typed (separately for fixed-format questions and for open-ended questions). Missing and invalid data rates were derived from each subject's responses.

RESULTS

Response Rate and Respondents

42 out of 49 people completed the survey for a response rate of 86%. 71% of the respondents were male. Two-thirds of the respondents had been employed at the lab for three years or less. 62% held a Ph.D.; 81% were working in a research-related position. The remaining 19% worked in administrative support. Respondents reported a mean level of general computing expertise of 8.0 (on a scale of 1 to 10, where 1 is "not at all competent" and 10 is "extremely competent"). Respondent characteristics were acceptably balanced across conditions.

Differences Between Faces and Text

The presentation condition had no effect on the subjects' reported satisfaction with their computing environment. Statistically significant differences between the groups did

Table 1

Attitudes and behavior as a function of presentation condition: text, neutral face, stern face. Unless noted, cells contain the mean rating on a 1-10 scale where 1 was the negative extreme and 10 was the positive extreme.

	Text (*n=15*)	Neutral Face (*n=15*)	Stern Face (*n=12*)
Attitudes to answering questions			
Were questions clear?	8.1	7.1	6.2
Were you comfortable?	8.7	8.2	8.3
Similar to face-to-face?	2.2	1.9	1.9
Want to continue?	4.6	3.4	3.0
Involvement			
Time spent (minutes)	24.1	26.8	35.0
Requests to repeat (freq)	N/A	6.5	8.0
Missing answers (freq)	3.6	3.3	2.7
Invalid answers (freq)	8.1	4.9	3.2
Unsolicited comments (words)[a]	57.1	24.9	57.2
Open-ended questions (words)[b]	65.6	57.8	114.9
Final comments (words)	30.0	37.9	68.3

[a] Comments made on fixed-format questions
[b] Excluding question about experience of answering questions

appear, however, for questions about the experiment and the experience of participating. See Table 1.

There were several main effects of presentation condition. Subjects in the face conditions rated the questions as significantly less clear ($F(2,39)=3.1, p=0.06$). Subjects in the three conditions differed significantly in how long they spent on the survey ($F(2,39)=3.86, p=0.03$) and in how much they wrote in open-ended responses ($F(2,39)=3.61, p=0.04$).

Subjects in the neutral face and text conditions typed the same number of words in answer to open-ended questions but people in the stern face condition typed almost twice as many words. They were quite explicit about their dislike of the face, with many suggestions for how to "improve" it. Probably as a result of hearing questions repeated and typing lengthy comments, they also spent significantly longer taking the survey than subjects in the other groups.

Subjects differed markedly in their assessments of the interviewer (see Table 2). Since text condition subjects had not seen any question asker, their responses tended towards the "don't know" range of the 10-point scale. Subjects in the face conditions expressed markedly more negative attitudes. Eight out of the eleven comparisons were ordered with the text condition most positive, the stern face condition most negative, and the neutral face condition in the middle. Six of the eleven comparisons were statistically

Table 2

Assessing the attributes of "the question asker". Cell values are means on a 10-point scale, where 1 is the most negative assessment and 10 is the most positive.

Attribute	Text *(n=15)*	Neutral Face *(n=15)*	Stern Face *(n=12)*	F-ratios *(dfs 2,32 to 2,37)*
Likable	6.5	3.4	3.1	8.98***
Friendly	5.8	3.9	2.8	4.19**
Intelligent	6.2	4.7	4.4	1.74
Comfortable	7.6	6.5	4.4	3.38**
Trustworthy	6.5	5.4	4.9	1.16
Cooperative	5.3	5.7	3.7	1.52
Weak	3.5	3.2	4.7	1.21
Stiff	5.8	8.0	7.9	2.65*
Happy	5.6	2.7	2.5	9.23***
Generous	4.1	2.9	2.8	1.54
Sad	3.1	5.1	5.7	3.83*

*=p<.1, **=p<.05, ***=p<.01

significant. Although it was not significant, subjects in the face conditions were less willing to "continue a conversation" than were the subjects in the text condition. There were no statistically significant differences in the rate of missing or invalid data, although the trend was towards greater accuracy in the face conditions.

Differences Between Expressions

We further explored some of the main effects by direct comparison of the two face conditions. Subjects who saw the stern face found the questions less clear and asked to have more of them repeated. People in the stern face condition spent longer answering the survey than did people in the neutral face condition (35 vs. 27 minutes; t=1.96; p=.06). They wrote more in response to open-ended questions about their computing environment (183 words vs. 96 words; t=2.0; p=.05). They also wrote more (but not significantly more) about the experience of answering questions (68 words vs. 38 words). There were no statistically significant differences in error rates or substantive answers and no difference in response to questions measuring question clarity, desire to continue answering questions this way, or subject comfort. There were no significant differences in assessment of the question asker, but the stern face was judged in a more negative light in *all* of the comparisons.

Differences Between People

We explored gender and education level of the subjects as possible covariates. The resulting cell sizes were very small and non-proportional, preventing formal analyses of variance. We did, however, see some suggestion of an interaction between gender and education level in subjects' reaction to the experience, with high education females providing the most negative ratings on 11 out of 15 scales.

DISCUSSION AND CONCLUSIONS

We demonstrated the feasibility of having people interact with a talking face for purposes of asking and answering questions. Given the untried nature of the talking face technology, it is encouraging that people performed similarly across face and text conditions in terms of substantive response to the actual topic of the questionnaire. Furthermore, faces did not lead to poorer performance through distracting people. On the contrary the performance of people in the face conditions was better than in the text condition as measured by (lower) number of invalid responses. Faces were more engaging than text as measured by how long people spent answering questions and in how much they wrote in response to open-ended questions. Further, the face with "more" expression (the stern face) led to greater engagement than did the face with "less" expression (the neutral face). Note that engagement does not mean liking. Respondents assessed the stern face less positively than the neutral face and assessed both faces less positively than the text condition.

Why would people spend more time, respond more, and respond more carefully to a face that they did not like? Social psychologists have found that the presence of another person usually serves to increase arousal and motivation on the part of someone asked to perform a task [35]. This can lead to improved performance if the task is not very complex or to degraded performance if the task is complex. The presence of another person apparently produces evaluation reminders and therefore leads people to try harder. We posit that the more expressive face (the stern face) in this study may have actually produced the most evaluation reminders of the three conditions. (One respondent said the face reminded him of a "stern school marm.") If so, these evaluation reminders would cause people to pay closer attention to their task than did the other two conditions.

Synthetic talking face technology will continue to improve. We will be able to manipulate expression more extensively and to expand the range of facial movements possible. We will be able to incorporate more natural eye movements and more expressive speech. The major question from an interface designer's point of view is how to use these capabilities.

This study raises many fundamental questions about the future nature of the human-computer relationship. If a computer is a social actor, what is its role (master, apprentice, partner?) and how does the face it wears affect the behavior and attitudes of the person interacting with it? How should the faces used for delivering information (for example, in kiosks) differ from the faces used for soliciting information (for example, in medical history applications)? Should human facial realism be a goal? If so, whose face should appear? Based on this preliminary work, we anticipate differences in the appropriateness of face-based interfaces depending on the social interaction demands of the situation, on the expression and gender of the faces used, and on individual temperament characteristics of the users.

The goal of HCI work with synthetic faces should not necessarily be to give a computer a human face but rather to determine when a face-like interface is appropriate. It is particularly important to follow up on the possibility of gender effects, particularly in light of HCI's commitment to interfaces that are equally accessible and acceptable to all intended users. Our results indicate that significant further research is necessary to identify the components of a satisfactory experience with a human-like computer interface.

ACKNOWLEDGMENTS

Thanks to Keith Waters, Tom Levergood, and Dave Wecker for invaluable technical collaboration.

REFERENCES

1. Apple Computer Inc. Knowledge Navigator. (Videotape)
2. Aronoff, J., B.A. Woike, & L.M. Hyman. Which are the Stimuli in Facial Displays of Anger and Happiness? Configurational Bases of Emotion Recognition. *Journal of Personality and Social Psychology*, 1992. **62**, p. 1050-1066.
3. Bailey, L., T. Moore, & B. Bailar. An Interviewer Variance Study for the Eight Impact Cities of the National Crime Survey. *Journal of the American Statistical Association*, 1987. **73**, p. 16-23.
4. Baker, M.J. & G.A. Churchill Jr. The Impact of Physically Attractive Models on Advertising Evaluations. *Journal of Marketing Research*, 1977. **14**(November), p. 538-555.
5. Binnick, Y.M., C.F. Westbury, & D. Servan-Schreiber. Case Histories and Shorter Communications. *Behavioral Research Therapy*, 1989. **27**, p. 303-306.
6. Bond, E.K. Perceptions of Form by the Human Infant. *Psychological Bulletin*, 1972. **77**, p. 225-245.
7. Bradburn, N. Response Effects. In *Handbook of survey research*, P.H. Rossi, J.D. Wright, & A.B. Anderson, (Ed). New York: Academic Press, 1983. p. 289-328.
8. Carey, S. Becoming a face expert. In *Processing the Facial Image*, V. Bruce, *et al.*, (Ed). Oxford: Clarendon Press, 1992.
9. Deiner, E., *et al.* Effects of Deindividuating Variables on Stealing by Halloween Trick-or-treaters. *Journal of Personality and Social Psychology*, 1976. **33**, p. 178-183.
10. Duffy, S.A. & D.B. Pisoni. Comprehension of Synthetic Speech Produced by Rule: A Review and Theoretical Interpretation. *Language and Speech*, 1992. **35**, p. 351-389.
11. Eichenwald, K. "Hi, voter. This is your president." in *New York Times*, New York. November 2, 1986, section 3: 19.
12. Ekman, P. (Ed.) *Emotion in the human face*. 2nd ed. Cambridge: Cambridge University Press, 1982.
13. Ekman, P., *et al.* Does Image Size Affect Judgments of the Face? *Journal of Nonverbal Behavior*, 1979. **4**(1), p. 57-61.
14. Gaver, W.W. Auditory Icons: Using Sound in Computer Interfaces. *Human-Computer Interaction*, 1986. **2**, p. 167-177.
15. Hess, U., A. Kappas, & K.R. Scherer. Multichannel Communication of Emotion: Synthetic Signal Production. In *Facets of emotion: Recent research*, K.R. Scherer, (Ed). Hillsdale, N.J.: Lawrence Erlbaum Associates, 1988.
16. Itou, K.S., S. Hayamizu, & H. Tanaka. Continuous Speech Recognition by Context-dependent Phonetic HMM and an Efficient Algorithm for Finding N-best Sentence Hypotheses. In *Proceedings of ICASSP*. IEEE Press, 1992.
17. Klatt, D.H. & L.C. Klatt. Analysis, Synthesis, and Perception of Voice Quality Variations Among Female and Male Talkers. *The Journal of the Acoustical Society of America*, 1990. **87**(2), p. 820-856.
18. Laurel, B. Interface Agents: Metaphors with Character. In *The art of human-computer interface design*, B. Laurel, (Ed). New York: Addison-Wesley, 1990.
19. Morton, J. & M.H. Johnson. CONSPEC and CONLERN: A Two-process Theory of Infant Face Recognition. *Psychological Review*, 1991. **98**(2), p. 164-181.
20. Nass, C. & J. Steuer. Voices, Boxes, and Sources of Messages: Computers and Social Actors. *Human Communication Research*, in press. .
21. Neidecker-Lutz, B.K. & R. Ulichney. Software Motion Pictures. *Digital Technical Journal*, 1993. **5**(2), p. 19-27.
22. Norman, D.A. & S.W. Draper (Ed.) *User-centered system design*. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1986.
23. Oren, T., *et al.* Guides: Characterizing the Interface. In *The art of human-computer interface design*, B. Laurel, (Ed): New York: Addison-Wesley, 1990. p. 367-381.
24. Reilly, W.S. & J. Bates. Building Emotional Agents. School of Computer Science, Carnegie Mellon University CMU-CS-92-143, 1992.
25. Resnick, P.V. & H.B. Lammers. The Influence of Self-esteem on Cognitive Responses to Machine-like versus Human-like Computer Feedback. *The Journal of Social Psychology*, 1985. **125**(6), p. 761-769.
26. Rich, C. Evaluating the Contribution of Discourse Theory to an Interactive System. In *Proceedings of AAAI Fall Symposium on Human-Computer Collaboration*. Raleigh, N.C., October, 1993.
27. Schuman, H. & S. Presser. *Questions and answers in attitude surveys: Experiments in question form, wording, and context*. New York: Academic Press, 1981.

28. Shneiderman, B. *Designing the user interface: Strategies for effective human-computer interaction.* Boston: Addison-Wesley, 1987.
29. Sidner, C. Using Discourse to Negotiate in Collaborative Activity: An Artificial Language. 1992, Workshop on Cooperation Among Heterogeneous Agents: NCAI Conference.
30. Takemura, H. & F. Kishino. Cooperative Work Environment Using Virtual Workspace. In *Proceedings of Computer Supported Cooperative Work.* Toronto, 1992.
31. Takeuchi, A. & K. Nagao. Communicative Facial Displays as a New Conversational Modality. In *Proceedings of INTERCHI '93.* Amsterdam, ACM Press, 1993.
32. Thorisson, K.R. Dialog Control in Social Interface Agents. In *Proceedings of INTERCHI '93.* Amsterdam, ACM Press, 1993.
33. Waters, K. A Muscle Model for Animating Three-dimensional Facial Expressions. *Computer Graphics,* 1987. **21**(4), p. 17-24.
34. Waters, K. & T.M. Levergood. DECface: An Automatic Lip-synchronization Algorithm for Synthetic Faces. Cambridge Research Laboratory, Digital Equipment Corporation Technical Report CRL 93/4, 1993.
35. Zajonc, R. Social Facilitation. *Science,* 1965. **149**, p. 269-274.
36. Zeithaml, V.A., A. Parasuraman, & L.L. Berry. Communication and Control Processes in the Delivery of Service Quality. *Journal of Marketing,* 1988. **52**, p. 35-48.
37. Zeithaml, V.A., A. Parasuraman, & L.L. Berry. *Delivering quality service: balancing customer perceptions and expectations.* New York: Free Press, 1990.