

Social Bandwidth : The Importance of Social Issues for Understanding the Effects of High-Bandwidth Interactions

Byron Reeves

Department of Communication

Stanford University

Reeves@leland.stanford.edu

Abstract

Panel III, Technology and Social Interaction, considered the implications of Internet2 for social perceptions and social interactions. The discussion below summarizes research opportunities that were discussed, focusing on the psychological possibilities and problems that could be included in a research agenda about the social effects of high-bandwidth technologies.

A basic premise of this consideration is that interactive media — including the new technologies considered at the Internet2 Sociotechnical Summit -- are not merely tools that enhance productivity and the efficiency of work. Rather, these technologies will influence social interaction, in some cases causing social issues to become more dominant than is true for present technology, and in other cases causing the success of new technology to be determined as much by social issues as by technical merit [19].

1. Introduction: Media experiences are social experiences

More engineering bandwidth means more social bandwidth. As technology can better communicate the breadth of information available in live human-to-human interactions, the natural tendency for people in mediated interactions will be increased attention to social cues — both as receivers and sources of information. Social issues can determine the success of interactions with less sophisticated technologies, including simple text, but an increase in social bandwidth should turn up the social volume even more. In real life, we commonly understand that social competence is

critical for success in commerce, education, politics, and interpersonal relationships. Technologies such as Internet2 should provide an infrastructure that assures this is no less true on the net.

2. Categories of social issues relevant to Internet2

There are at least three ways that social issues might be categorized as Internet2 research is planned. First, the enhanced bandwidth of Internet2 will enable more socially complete exchanges between people. More social cues will be available to send and receive as people use technology to facilitate contact with other people. Second, more bandwidth will enable technology that automates social interaction such that computers can become social actors themselves. Computers will become less distinguishable from real social partners in relationships between users and machines. Third, new Internet technologies will present different social opportunities for access to new forms of interaction, with potentially negative effects accruing to those who cannot participate.

These categories are summarized below, including example questions about the specific technologies that have the highest likelihood of enhancing social interaction. I follow this with a suggestion for the establishment of an archive with data about social transactions and new technology. At the end, there are comments about what we might learn from the introduction of other technologies during this century as they evolved from the promises of their inventors to the real worlds of the marketplace, classrooms and living rooms.

2.1. Social Bandwidth will enhance computer-mediated communication

More bandwidth will allow computers to offer more faithful representations of the people with whom we interact. This is often categorized as computer-mediated communication (CMC), and it is the simple value proposition of most conferencing technologies including those proposed for scientific collaborations, teacher-student interactions, enterprise-client transactions, and doctor-patient relationships. Computers, because they will be able to exchange more information more quickly via Internet2, will diminish social distance while preserving rich social presentations.

All the cues available in face-to-face encounters -- those used on purpose as well as those that are presented incidentally — will be more faithfully represented in high-bandwidth exchanges. This includes facial expressions, gestures, posture, motion, and even touch, texture, smell, and size. In addition, the exchanges may carry social information that is not possible in the *real* world but that is socially consequential; for example, the ability of interactants to view themselves as well as their conversation partners. All of this information can determine a range of social responses to the interactions including personality attributions, emotional experience, social memory and attention.

It will not be the case, however, that social dimensions of interaction will be added in a single increment, and even if that were possible, each possible feature of social interaction would not be equally useful in this type of communication. Consequently, research should try to find out which pieces of social information carry the biggest bang per byte for different interaction criteria (e.g. commerce versus education versus entertainment).

Here are example elements of social presentations — all requiring radically different amounts of bandwidth -- which could be studied for their potential to augment social responses in mediated interactions:

2.1.1. Image fidelity

What level of image resolution is necessary for effective representation of interactants? There is evidence that the value of high-definition images may be overestimated [16, 19], especially relative to the importance of audio fidelity. There also exists research about the importance of examining perceptual systems as they relate to

the processing of different picture formats [5, 7]. Fidelity should be studied before assumptions are made about the universal benefits of high-definition images.

Is there a minimum threshold of fidelity necessary to engage users? Do higher resolutions that increase a sense of presence also increase arousal? There is good evidence that the structural features of mediated presentations are as effective, and sometimes more influential, than picture content as a determinant of excitement when viewing advanced displays [9]. The consequences of higher arousal images are important; people are given more attention to them, and they are better-remembered [19].

How might the effects of image fidelity change for different types of interactions? Would higher fidelity educational presentations be more compelling? Would they enhance learning? Or could visual presentations become so compelling that they draw attention away from the presentation of information? These questions center on how new media, and especially their compelling visual features, compete for the same mental resources that also enable learning [9].

How does image fidelity influence the perception of faces, a critical component of social perception (e.g. more facial cues are available as resolution increases)? Will interactants become more focused on the presence of other people and less able to concentrate on what people say? The perception of faces, especially in the context of new media, has generated a substantial literature already [10]. A major point in this literature is that people are highly sensitive to minute facial features. Changing the way that these features are presented can change the details, and even the essence, of a mediated exchange.

2.1.2. Image size

There are several choices in formats for displaying information. How might display size influence social interactions? Do the personalities of interactants change? Can larger displays enhance learning by providing more places for information to exist?

There is evidence that display sizes do effect how people respond to media [19]. Larger screens are more arousing [3]. People on larger screens have personalities that are more dominant, compared to responses to the same images on smaller devices [23]. As the number of different sizes of these devices increase, these differences may become even more important. And as display sizes increase, a likely by-product

of more bandwidth, information presentations will become more arousing.

There are important questions left to be studied, as people are able to access similar information in different formats. What display technologies -- Personal Digital Assistants (PDAs) to wall screens -- are most appropriate for different types of Internet2 interactions? What types of interactions (and which social responses) benefit the most from small or large displays?

2.1.3. Timing and delays

Timing is socially meaningful. Even changes in seconds can influence how people understand different comments in an interaction. Media impose both time delays (e.g. download delays that are longer than real life counterparts) and time enhancements (e.g. looking through large databases in a tiny fraction of the time it would take in real life). Very little is known, however, about how these changes in the time course of mediated interactions change the perception of information.

2.1.4. Information synchrony

At present, there are significant problems in the synchronization of information presented in lower bandwidth environments. In teleconferencing systems, for example, it is common that lips are not synchronized with speech, or gestures not synchronized with language. There is already evidence that these problems can cause significant changes in how people understand the information that is exchanged, and how they evaluate the technology that enables the exchanges [19].

The social consequences of information that is not synchronized (e.g. mouths with speech, gestures with vocalizations) can spell disaster. When asynchrony occurs, people are perceived as deficient, unworthy of positive regard. While this is an unconscious response, it does indicate the difficulty in expecting that people will be able to somehow discount technical problems. Rather than blame the technology, people blame the humans at the other end. Technical problems become perceptions of social awkwardness. Documenting how higher bandwidth exchanges can eliminate these problems will offer significant positive evidence to support wider use.

2.1.5. Display of self-knowledge

High bandwidth exchanges will make available, in real time, a greater range of

information about an interaction. People will not only see the other people joining a conversation; they will often also see themselves. What are the social implications of giving interactants information about themselves? (e.g. What happens when a speaker can see him or herself on a monitor while talking?).

The first studies of the effects of this additional information [13] show that pictures of oneself are often evaluated as separate social actors. Also, pictures of oneself increase self-consciousness and can alter interaction styles. How will more complete access to this information change, discourage or encourage use of advanced systems?

2.1.6. Multiple modalities

Increased bandwidth almost surely brings the possibilities of interacting with other people via different social channels. There are several important questions to ask about the redundancies in these channels, the advisability of using all channels that are available, and the social consequences of having asymmetry in the channels available to two different interactants.

Redundancy in expression can be valuable in communication, and there is already a significant literature that shows when and under what conditions redundancy is useful. Most past research, however, deals with the presentation of information to a recipient, not the interaction between two people who are *both* sources and receivers of information. Much effort will go into enabling multiple-modality interactions, and there is already evidence to support the value of these exchanges [17].

There is also interesting research that examines interactive exchanges where there is asymmetry in the channels available during the interaction [14]. It can be disconcerting, for example, when one person can be seen and heard, and the other only heard. As more possible channels are added to interactions as bandwidth increases, research should further examine the consequences of uneven access to different modalities.

2.2. Social Bandwidth will enable compelling automated social relationships

More bandwidth will allow computers and their interfaces to become plausible social actors themselves. This has the possibility to dramatically move human-computer interactions (HCI) away from cognitive manipulations (e.g. pressing interface buttons, arranging information hierarchies) to more social tasks (e.g. conversing

with an automated agent). The implications of enhanced social intelligence in this domain may suggest even more problems and opportunities than in the case of CMC, mostly because the questions here relate to the promise of artificial life.

The discussion of automated social relationships is often relegated to virtual worlds, and indeed there are many important questions about technology at that level of sophistication (see Panel II on Teleimmersion, Embodiment and Interface Design). Sophisticated social responses occur, however, regardless of the sophistication of the technology that presents the relationships. People cry when they read, and can easily be bored in a virtual world. In fact, it may be the lesser social cues that surface first as Internet2 is deployed. While the virtual worlds are readied for the future, new found bandwidth is sure to be used for less sophisticated — but socially critical — enhancements like speeding up interactions or allowing computers to adapt to the personal characteristics of users

Many of the potential research questions in this category should be similar to those for CMC. From the perspective of technology users, automated interactions are similar to person-to-person exchanges even though they replace the other person with a machine. Other issues, however, will be unique to automated agents. Here are some examples:

2.2.1. Representation

When social actors are represented in automated transactions, they can appear in many different formats (e.g. photorealistic moving pictures, animated characters, and plain text). Each level of sophistication in this representational hierarchy suggests different competencies of the social actor, and different ways that the intelligence of the computer controlling the agents can be understood.

How will different forms of agent representation influence automated interactions? Will different representations change perceptions of credibility, intelligence, and believability? Will people expect more intelligence from an agent if they are represented with a high degree of realism? At what level of consciousness are automated agents processed (e.g. are people able to constantly consider the artificiality of representations or do they attribute realism automatically when cues signal that agents are close enough to the real thing)?

There are already several preliminary studies that suggest differences in perceptions of

interactions based on the type of representations used [e.g. 2]. This is a significant area for new research; however, given the extremely large number of possible representational styles that high-bandwidth computing will allow, many of the possible presentation formats have not been studied.

2.2.2. Recognizing speech

One of the most significant increments to social presence in high-bandwidth computing is the presence of human speech. Faster networks will make it possible to run sophisticated speech recognition software in real time, and make it possible to perform human-quality text-to-speech synthesis. Preliminary research has already shown that speech features of interactive media are highly social, and their presence in automated interaction increases the social feel of an exchange [12].

Much is left to be studied, however, about how people will respond to systems that recognize speech. Is the ability to speak to an automated agent different than the ability to write text or even to speak to a machine when there is no explicit agent present? When people are allowed to speak to a computer, as opposed to using a keyboard or mouse, they are more likely to respond to the machine as if it were another person.

There are also social questions that could influence the success of speech recognition as a technology. Currently, many speech recognition systems depend on the ability of interaction designers to anticipate the range of utterances that people will use in particular parts of an interaction. These determinations, accomplished with the help of linguistic analyses, are the basis for constructing recognition grammars that specify a subset of utterances (out of infinite possibilities) that are most likely. The social scientific aspects of these constructions, however, are unstudied. How can exchanges with automated systems be constrained so that people will be more likely to make in-grammar utterances?

2.2.3. Effects of artificial qualities of interactions

A solid premise in the construction of automated systems is that successful features of human-human exchanges should be transferred, as completely as possible, to human-computer interactions. There are several cues in automated exchanges, however, that are artificial, and little

is know about how they will change the nature of the interactions.

One of the most significant artificial cues is TTS (text to speech) translation. With present technology, automated TTS voices, while comprehensible, do not sound like plausible human voices. To date, most of the research about TTS has focused on the issue of whether people can understand the words spoken. Social responses to TTS is less well understood, although there are preliminary studies that show TTS voices that vary on characteristics like mood, gender and personality have effects similar to real voices that vary on the same qualities [12]. There are also open questions about social responses when familiar and unfamiliar automated features are combined (e.g. what happens when recorded audio — the familiar sounds of recorded speech -- is mixed with speech synthesis --an unfamiliar voice)? This is a concern for many systems that currently mix recorded audio with TTS translations for content in e-mail and other presentations.

2.2.4. Automating dialogue

The construction of automated social dialogues for high-bandwidth systems is a complicated venture. A design goal in these systems is to accent the non-linear possibilities in the interactions, a prospect that dramatically complicates the design process. There are studies, mostly from an engineering perspective, that address issues surrounding the creation of tools that allow designers to construct automated dialogues [1]. There are no studies, however, that address the influence of dialogue design tools on the people who actually compose the interactions.

Automated dialogues will complicate the difficult collaboration between engineers and artists. What kind of tools can be created that minimize the need for designers to write code, and maximize the flexibility of designers to explore non-linear scripts with the same ease that they edit natural language?

2.2.5. Adaptivity

One of the primary motivations for constructing automated dialogues is the possibility for computers to adapt to people who use the systems — both within and between interactions. A computer, more easily than a human, can keep track of interactions, recall user characteristics, and bring this information to bear during interactions. This type of adaptation is a

primary characteristic of successful human-human dialogue.

There is preliminary evidence that adaptation is a positive characteristic of automated transactions as well [19]. People enjoy relationships with technology more when systems adapt to their needs, styles, time as in times or schedule, and place. Users think better of themselves and their own work during the interaction, and they are willing to stick with the interactions longer when automated systems show some sign of knowing who they are.

Much is left to study, however. How much adaptation is necessary to convey a sense of personalized interaction? What are the consequences of too much adaptation? Is adaptivity more useful in certain social domains (e.g. would it be more useful to recognize and respond to emotions, personality, social roles, or moments of frustration)?

2.2.6. Emotions, personality and social roles

Automated social actors will be perceived to have many of the same features that determine social responses to real people. There is good evidence that computer agents are thought to have emotions, personality and to occupy different social roles (e.g. gender and occupational roles) [19]. These social attributions to automated agents should be even stronger when there is more bandwidth in which to develop and present social information.

Research should continue to evaluate the effects of these social roles on people's acceptance of and responses to automated agents. Ignoring these roles could complicate transactions. We know already, for example, that people would prefer to work with computer agents that match their own personality [19], people respond emotionally to the emotional expressions of agents, and that people respond differently to male agents than they do to female agents that perform the exact same functions [15].

2.2.7. Errors

Automated systems will always make errors. These could be anything from errors in speech recognition to saying the wrong thing at the wrong time in an automated dialogue. It is likely that higher social bandwidth in an interaction will offer more chances for social errors of consequence. It is not necessary, however, that deployment of automated systems wait until accuracy is near perfect. Humans make errors as well. The difference, however, is that humans

are experts at repairing them *during* an interaction.

Very little is known about the possibilities for high social bandwidth interactions to allow for social errors to be corrected. Research could determine the most effective strategies in human-human discourse, and test their application to automated systems. Can following simple social rules for apologies, explanations and disclosures lead to better perceptions of systems and greater satisfaction and learning?

2.2.8. Evolution over time

Many of the proposed automated social interactions will persist over long periods of time, perhaps months and years in the case of health systems, financial tracking systems and the like. There is no research available yet about the long-term consequences of these relationships. How should these relationships change and mature over time? Will people tire of automation over time and avoid automated systems? What forms of adaptivity, over longer time periods, will help relationships mature?

2.2.9. Social displacement

There is disagreement about whether participation in automated relationships, especially those that displace live relationships, will cause changes in how people relate to other humans [22]. One possible effect of concern is that people will seek fewer live interactions because of a preponderance of automated interactions.

This is already a documented effect in the area of political communication [11]. People who spend a lot of time with mediated political figures (on television and in newspapers) are less likely to engage their friends and family in discussion about politics. It would be useful to study similar effects in other areas of social interaction, for example, on-line education, commerce, and religious participation.

2.2.10. The ethics of well done social simulations

There have already been several significant reviews of the ethical considerations when machines can more fully simulate *real* social interactions [4]. While ethical concerns go beyond psychological and sociological research issues, they are critical components to a full appreciation of the effects of these systems.

Several questions should continue to be addressed. What type of information should automated systems have to provide to users

about how they work? Should these systems have to say who is controlling them? What guarantees about privacy should be given and when during a relationship?

2.3. Who has and needs access to increased Social Bandwidth?

The third set of social questions, different from the first two, asks about the people who will participate in the interactions. There are two different ways in which the demography of these social relationships may be important. The first set of questions concerns access. Regardless of the nature of positive or negative consequences that result with increased social bandwidth, it is almost certain -- as has been true with all other media technology this century -- that some people will be barred from participation.

Here are some sample questions with respect to access:

2.3.1. Economics of access

Questions about economics as a determinant of access are familiar in considerations of all media. What will be the economic barriers to participation? Will some groups be left out and with what consequence?

There are some arguments that bandwidth will be plentiful enough that connection time will be essentially free, but even if that is true, there are still larger economic questions about who will be able to afford and maintain the appliances needed to make the connections [6]. There is already research on Internet access that shows considerable differences in the ability of different economic groups, and on different demographic groups like race that are correlated with economic resources [8]. These studies demonstrate that there is good reason to doubt that higher bandwidth services will fare any differently and, if newer services provide even greater value, access limitations could even become more important.

2.3.2. Physical barriers to access

All construction projects benefit if access for disabled persons is considered from the beginning, and advanced Internet construction is no different. The opportunities to allow access to all will be greater with higher bandwidth services because redundancies in information presentation will not be limited by constraints on speed of data transmission.

There are several important projects underway that study and build alternative access to

technology [18]. These efforts should continue and should focus in particular on the most useful technologies that will allow universal access. Examples of the most promising of these technologies include eye-tracking navigation systems, voice-recognition systems for dictation, and auditory-information systems that complement visual presentations.

A second sense in which demography is important concerns how groups of people will be affected differently by enhanced social interactions. Will the problems and opportunities fall to some more than others? Some sample questions include:

2.3.3. Children

With the introduction of every new medium this century, children have been usefully studied as a special audience. Children process media differently than adults and are susceptible to a greater range of potentially harmful effects. Also, children spend considerable more time with media than adults, and we already know that this is true for many types of current Internet offerings [21].

One important research orientation when studying newer Internet technologies is to compare the findings about children and television with those for newer technologies. The literature about children and media is one of the most substantial among all areas of communication. And there are many lessons from that research that have been shown to recur as each new medium was introduced this century [24]. Research often moves, for example, from concerns about physical effects to concerns about psychological changes. Will the same progression of research be true for Internet technologies?

It is also true that children are the most prone to consider media presentations as *real*. This is true even for standard display formats, like small television sets, which are not particularly compelling. Studies about newer technologies should ask whether advanced displays and engaging interactions make children even more likely to accept mediated information as real. Might these compelling displays actually limit the ability of educational presentations to be effective? When interactions become too compelling, children may be more likely than adults to focus on the displays themselves, ignoring the lessons within.

There are also more reasons to consider physical health risks associated with advanced technologies because the presentations

themselves involve more complete physical and perceptual participation. Are there health risks to children like motion sickness, eyestrain, and disorientation?

Compelling high-bandwidth interactions may also offer health benefits. Can automated interactions enhance motor coordination, or offer effective treatments for mental and physical disorders through the use of automated agents that interact with children about health issues, and do so with regularity and in an entertaining context?

2.3.4. Education and expertise related to technology

It is tempting to relegate concerns about technology effects to those who do not understand technology. The social responses considered here, however, have been shown to be common for all people, a result that has appeared consistently across a range of research [19]. Computer scientists, just as much as first-time users, for example, treat computers like real social actors in a conversation, and they are equally affected by the different ways that technology presents other people in interactions.

The fact that social responses occur for everyone suggests that many research questions single out the most sophisticated users. How do higher social bandwidth exchanges change the social relationships between researchers who collaborate from a distance? How are relationships changed between doctors and patients, students and teachers, or competitors in different laboratories? Many of these relationships in real life succeed or fail as much based on social intelligence as intellectual merit. Will Internet2 exchanges accent the social, relative to current exchanges that may offer less social information?

2.3.5. Personality differences

There are some personality differences which manifest themselves in the use of interactive media. There is evidence, for example, that shy people are less likely to engage in mediated interactions, and use technology differently when they do. There is counter evidence, however, that suggests that shy people become less so when enabled by technology [22]. Personality differences in willingness to use technology and the effects of use on people with different social orientations are useful directions for future research.

Another individual difference that has relevance to the social responses to Internet

technologies is the degree to which interactants have an external or internal locus of control. Those who are externally oriented think that *others* are in control, and these people are more likely to respond favorably to social cues in interactive media [20]. Those who are internally focused think that success with technology depends on their *own* skills, not other people's intervention, and they are less tolerant of social features in mediated interactions (e.g. animated computer agents). This work suggests future studies about how to build interfaces that allow people to choose a degree of social facilitation in interfaces that best suits their orientation. This could have implications for constructing appropriate on-line educational experiences as well as facilitating commerce transactions.

2.4. Data about social interactions using technology

There is an important opportunity for social research that relates to all of the questions addressed above. Almost every question about high-bandwidth social exchanges requires information about the actual transactions that people have with technology. What do people do with technology? How are various responses related to different stimuli and opportunities in the transactions? How does the nature of the transactions change over time?

Information relevant to these exchanges can be captured in the buttons and choices that people make as they use the technologies. Most of the research questions would benefit if there were an effort to establish a library or repository of data about such transactions. Because computers manage these transactions, it is relatively easy (certainly easier than with other recording devices in a laboratory) to record transactions for later analysis. This should be a high priority for research funding. It will ensure that there is useful data to answer many of questions that emerge in the future.

3. What should social mean?

It is difficult to open a psychology textbook, whether about social, cognitive, perceptual or physiological issues — and not find that almost every chapter relates to this inquiry. This is an excellent reason to view the lists of sample questions above as only a beginning. There is every reason to think that this list can be lengthened (perhaps more so than the lists for

other media of this century). Computers can do a lot of social things.

A fundamental reason for the breadth of questions about Internet2 is the breadth of experiences afforded by the technology. Many of these experiences are actually quite familiar, at least psychologically. Granted that the experiences are produced using methods that are technically novel, even mysterious. But they more often than not bring users into interactions that share underlying similarities with real world encounters. Consequently, there seems little reason to limit the term *social* as it might apply to interactive technology. All psychological concepts seem like fair game — personality, emotions, social roles, memory, attention, perception, social norms, etc.

It is also important to look toward definitions of *social* that are not constrained by information content. It is unlikely that Internet2 will be limited to commerce, education, politics, advertising, public affairs, music or any other content category. Rather, the term *social* will apply to all of these domains, as has been the case with other media.

This breadth may complicate rather than simplify the categorization of research. It will mean we need both a large agenda and a lot of different people, those who know about the technology and those with primary expertise in the social sciences themselves. But in any case, the breadth is important. This is the only way to insure that technologies like Internet2 are defined as social phenomena in addition to their technical specifications.

4. The relevance of social issues: lessons from the study of "old" media when they were new

It may be tempting to consider social questions about Internet2 as premature. After all, those who are planning Internet2 mostly tout opportunities for scientific and educational exchange. That certainly doesn't sound extraordinarily social. Bring on the engineers instead.

An important history lesson suggests that ignoring social issues would be folly. Every important new media invention this century has been initially energized by similar promises. In reality, however, the social domain catches on fast. For television, film and radio, we progressed from the promises of education, travel, and a larger public sphere to content

largely driven by entertainment and compelling, famous, personalities. For the Internet, the progression has been far quicker. In 1994, the Internet was the domain of physicists trading data. Only five years later, we can hardly remember those first academic uses. The top Internet sites, and a huge majority of the traffic, are about commerce and entertainment.

An important implication of this history concerns at what point we decide to emphasize social issues in research. The history lesson is that sooner is better. We will know more about how the technologies will ultimately be used by considering social responses at the outset. This is often extremely difficult, in large part, because it's difficult to get people to focus on social issues at the same time as technical wonders. We'll know more in the long run, however, by starting social research at the outset. And this knowledge should benefit those whose business depends on the technology as much as to those whose interest is primarily in social research.

References

- [1] R. Agarwal, "Towards A PURE Spoken Dialogue System For Information Access," *Proceedings Of the ACL/EACL Workshop On Interactive Spoken Dialog Systems*, Madrid, Spain, 1997, pp. 90-97.
- [2] J. Cassell, J. S. Sullivan, and S. Prevost, *Embodied Conversational Agents*, MIT Press, Cambridge, MA, (in press).
- [3] B. Detenber and B. Reeves, "A Bio-Informational Theory Of Emotion: Motion And Image Size Effects On Viewers," *Journal of Communication* 46, 1996, pp. 66-84.
- [4] B. Friedman, *Human Values and the Design of Computer Technology*, Cambridge University Press, New York, NY, 1997.
- [5] T. Fukuda, "New Electronic Media And The Human Interface," *Ergonomics* 33 (6), 1990, pp. 687-706.
- [6] S. Greenstein, and M. Lizardo, "Determinants of the Regional Distribution of Information Technology Infrastructure in the U.S.," eds. D. Orr and T. Wilson, *The Electronic Village: Public Policy Issues of the Information Economy*, C. D. Howe Institute, Toronto, Canada, 1999.
- [7] J.E. Hochberg, "Representations Of Motion and Space In Video And Cinematic Displays," eds. K.R. Boff, L. Kaufman and J.P. Thomas, *Handbook Of Perception And Human Performance*, John Wiley and Sons, New York, NY, 1986.
- [8] D.L. Hoffman and T.P. Novak, "Bridging The Racial Divide on the Internet," *Science* 280, April 17, 1988, pp. 390-391.
- [9] A. Lang, R.F. Potter and P.D. Bolls, "Something For Nothing: Is Visual Encoding Automatic?" *Media Psychology* 1 (2), 1999, pp.145-164.
- [10] D.W. Masaro, *Perceiving Talking Faces: From Speech Perception to the Behavioral Principle*, MIT Press, Cambridge, MA, 1998.
- [11] D. Mutz, *Impersonal Influence*, Cambridge University Press, New York, NY, 1998.
- [12] C. Nass and L. Gong, "Maximized Modality or Constrained Consistency?" *Proceedings of the AVSP 99 Conference*, Santa Cruz, CA, 1999.
- [13] C. Nass, E-Y. Kim and E-J. Lee, "When Your Face Is The Interface: An Experimental Comparison Of Interacting With One's Own Face or Someone Else's Face," *Proceedings of the CHI Conference*, Los Angeles, CA, 1998, pp. 148-154.
- [14] C. Nass and Y. Moon, "Machines and Mindlessness: Social Responses to Computers." *Journal of Social Issues*, (in press).
- [15] C. Nass, Y. Moon and N. Green, "Are Computers Gender-Neutral? Gender Stereotypic Responses to Computers," *Journal of Applied Social Psychology* 27 (10), 1997, pp. 864-876.
- [16] W.R. Neuman, A.C. Crigler and V.M. Bove, "Television Sound And Viewer Perceptions," *Proceedings of the Joint IEEE/Audio Engineering Society*, February 1991.
- [17] S.L. Oviatt, P.R. Cohen, M.W. Fong and M.P. Frank, "User-Centered Design of Spoken Language and Multimodal Interfaces". *IEEE Multimedia* 3, Winter 1996.
- [18] J. Perry, E. Macken, N. Scott, and J. Mc Kinley, "Disability, Inability, and Cyberspace," ed. B. Friedman, *Human Values and the Design of Computer Technology*, Cambridge University Press, New York, NY, 1997, pp. 65-89.
- [19] B. Reeves and C. Nass, *The Media Equation: How People Treat Computers, Television, and New Media Like Real People and Places*, Cambridge University Press/CSLI, New York, NY, 1996.
- [20] R. Rickenberg and B. Reeves, "The Effects of Animated Characters on Anxiety, Task Performance and Evaluations of User Interfaces." (paper submitted for presentation to *CHI 2000*, 1999.)
- [21] D.F. Roberts, U. Foehr, V. Rideout and M. Brodie, *Kids and Media at the New Millennium*, Kaiser Family Foundation: Menlo Park, CA, (in press).
- [22] S. Turkle, *Life On The Screen: Identity In The Age Of The Internet*, Simon & Schuster, New York, NY, 1995.
- [23] D. Voelker, *The Effects of Image Size and Voice Volume on the Evaluation of Represented Faces*, Doctoral Dissertation, Stanford University, 1994.
- [24] E. Wartella and B. Reeves, "Recurring Issues in Research on Children and Media," *Educational Technology* 23, 1983, pp. 5-8.