

Being real

Judith S. Donath

1.0 Introduction

How we know each other – how we perceive and construct the identity of our fellow humans – is a difficult question, entangled in the subjectivity of our social perceptions and the many and often opaque motivations of those whom we are attempting to comprehend. It is a difficult question in the real world, where signals as subtle as the slightest raise of an eyebrow can indicate, to those astute enough to notice, a wealth of information about one’s allegiances and beliefs – and where we exist amidst a cacophonous abundance of such signals. It is an even more difficult question in the virtual world, where the medium has nearly silenced the cacophony, leaving us to seek scarce hints of identity amidst the typed messages and static, stilted homepages.

This chapter will address the problem of teleidentity: how do we – or do we – “know” another person whom we have encountered in a mediated environment? The epistemological ramifications of this question are several. One of the most interesting and significant is the issue of credibility: how do we know whether or not to believe what we are told by someone? The traditional philosophic approach holds that sincerity and competence are the underpinnings of credibility [AUDI 1998]; in the mediated world, not only is our judgement of these matters made more difficult by the sparsity of social cues, but the very issue of the speaker’s identity, generally taken for granted in the physical world, becomes a source of doubt and an item requiring its own adjudication of belief and justification. There are also ethical ramifications. Knowing something about a person’s social identity is fundamental for knowing how to act toward them, for the complex rules of social conduct that govern our behavior toward each other cannot function in the absence of information about the other [HOLLAND AND SKINNER 1987]. The philosophical ramifications of teleidentity are of more than theoretical interest. The online world is growing in both size and significance: it has become a place that people go to for medical advice, investment strategies, news of the world; it is a place people turn to for community and support. We need to know something about the identity of those who supply information in order to assess its veracity and of those with whom we socialize in order to build a functioning community.

This essay approaches these issues by focusing on a question with special resonance for both technologists and philosophers: can one tell if the person at the other end of an online discussion is indeed a person? The problem of “other minds”, while of perennial philosophical interest, is not one that normally intrudes upon everyday life. One concludes either that others do indeed have minds (the pragmatic approach) or that the state of others’ minds is unknowable (the skeptical approach) and then goes about one’s daily business. The advent of computer-mediated communication – and, particularly, the advent of communication between man and machine – has changed this dramatically. Suddenly the question of other minds, as in “is the being with whom I am speaking in any way conscious or intelligent?” is no longer a rhetorical question asked only in ironic exasperation, but a pressing problem addressed with increasing frequency by ordinary people (i.e. non-philosophers).

And it is a question with a new twist: simple programs, about whose lack of intelligence there is little debate, have been shown to be capable of producing humanlike responses within the limited domain of textual conversation, convincingly enough to beguile some people into thinking they were indeed conversing with a fellow human. Epistemological debates about the existence and knowability of “other minds” often pose a skeptical view hypothesizing that the other person may actually be a robot or other non-conscious being. The mediated computational environment makes this a very real possibility. Increasingly, computer programs are participating in mediated conversa-

tions. Thus far their human partners have, with a few exceptions, soon recognized the machine essence of their ostensible companion. However, it is likely to be only a matter of time until a machine can truly masquerade as a human, at least via some media. Most writing on this topic is concerned with the design and nature of the “intelligent” program; here, we are interested in a different aspect of this problem – in people’s perception of these programs. How do people distinguish between a human and a machine claiming to be human?

2.0 Recognizing the human

ELIZA

ELIZA: Hello, I'm ELIZA

Me: I'm trying to write a paper

ELIZA: Does it please you to believe I am trying to write a paper?

Me: No, not really

ELIZA: You are being a bit negative

HEX

Me: Hello

Hex: Hello. Now, that's the formalities over!

Me: How are you?

Hex: I was feeling great until you started acting silly!

Me: I'm trying to write a paper

Hex: What a coincidence, you are trying to write a paper - that makes two of us!

Me: So's ELIZA

Hex: I have played around with ELIZA a bit. It's pretty obvious that it is a computer these days, but apparently it fooled people into thinking it was human when it was first written.

Me: Do you want to write my paper?

Hex: Hehehe - that's supposed to be a *question*??? So, what do you do for a living?

These are snippets of real conversations. Although the conversations were real, two of the participants were not – at least not in the sense of being human. ELIZA and Hex are programs (ELIZA ND; HUTCHINS ND). They are simple programs, essentially just linguistic parsers, with no underlying intelligence. Yet we easily attribute intelligence, humanity, and even personality to them: ELIZA seems distant and oddly disengaged; Hex seems louder, a bit obnoxious and rambunctious.

ELIZA was written in the early 1960s by MIT professor Joseph Weizenbaum in response to Alan Turing’s proposal of the “Imitation Game” (now commonly referred to as the Turing Test) as a test of whether a machine is intelligent. In his 1950 paper “Computing Machinery and Intelligence” Turing posited that while the question “can a machine think” is of great philosophical interest, it is too vague and untestable to be meaningful. Instead, he said it can be usefully substituted for by a game in which a human judge, faced with both a human and a computer trying to pass as human, tries to determine which one is the human. The test, Turing suggested, should be conducted via teletype, thus limiting the zone of inquiry to conversational and cognitive abilities. Turing predicted that by around the year 2000, computers would “win” about 70 percent of the time.

Turing based his “test” on a parlor game in which a man and a woman, hidden behind a curtain and both professing to be female, communicate by notes with another player who attempts to figure out which one is actually a woman. The game does not test one’s ability to *be* another gender, it tests mastery of the *knowledge* that goes into performing the gender role. This distinction hinges on mediation: if the communication were not mediated – if the judge could see and hear the contestants directly – playing the deceptive role would be vastly more difficult, involving physical transformation as well as knowledge and role-playing. By making the communication mediated, limited only to written notes, the ordinarily easy task of telling male from female becomes difficult.

Turing's paper has been interpreted many ways, ranging from a manifesto proclaiming the ease of near-term intelligent machinery to a statement of extreme skepticism highlighting the unknowability of all aspects of other minds. Whether Turing believed that a machine that could pass as human had to be able to think – or might possibly be able to think – is unclear. He devoted a considerable amount of the paper to refuting arguments stating machines cannot think, in a manner that suggests that he thought they might well be able to. He calls the Imitation Game a “more accurate” form of the question “Can machines think?”, implying that he believed there was some essential connection between acting like one was thinking and actually thinking. Yet the explicit parallel he draws between the gender-based parlor game (in which the difference between imitating a woman and being a woman is clear) and the computer/human test suggests that his primary concern was functional: he was interested in whether the computer could act like a human, rather than in the deeper, but unknowable question of whether the computer was in essence like a human. And finally, he says that the issue also hinges on semantics:

The original question, “Can machines think?” I believe to be too meaningless to deserve discussion. Nevertheless I believe that at the end of the century the use of words and general educated opinion will have altered so much that one will be able to speak of machines thinking without expecting to be contradicted.

Anthropomorphism is not new and throughout history many phenomena have been accorded human characteristics. Today, machines are indeed commonly referred to as if they were conscious: Reeves and Nass have shown that only minimally humanlike behaviors are need to trigger a social response [REEVES AND NASS 1996]. Yet our relationship to the anthropomorphized machine is complex: when asked directly whether a computer can think, many would say “no”, although in actuality they interact with the machine as if it were a thinking being, attributing volition to it and reacting its “opinions” much as they might to another person's.

When Joseph Weizenbaum created ELIZA his goal was certainly not to create a program that would fool people into thinking it was human. Rather, he hoped to show that a program that could parse natural language and had some simple heuristics for formulating responses – a program with no pretence of intelligence – could play the Imitation Game reasonably well. His intent was to demonstrate that this game was not a “more accurate” test for intelligence since patently unintelligent machines could be made to respond in a believably humanlike way. Much to his dismay, many people met ELIZA with great enthusiasm, embracing it as an intelligent conversational partner; some even suggested that ELIZA-like programs could replace human psychotherapists. These responses greatly discouraged Weizenbaum, who effectively retired from AI and became a crusader for humanism in the face of advancing technology.

People's enthusiasm for ELIZA is at first glance surprising. She (it?) responds by rephrasing your words back as a question or a general query about your thoughts and feelings; the effect is chilly and stilted. Why did people become so involved in talking to her? One factor is that Weizenbaum introduced ELIZA as a Rogerian psychotherapist, whose method is to reflect patients' questions back to them to elicit further communication. This scenario gave people a context in which ELIZA's behavior seemed reasonable and rational – almost human.

As conversational programs go, ELIZA is quite primitive and few people who interact with ELIZA are actually fooled into thinking she is human. More sophisticated systems have, however, been known to converse undetected for a considerable time. A yearly contest, the Loebner Prize competition, offers \$100,000 to the first program that can pass a fairly rigorous version of the Turing test [LOEBNER 1999]. Although the prize money remains unclaimed, many of the programs have fooled some of the judges for some of the time, holding their own in discussions about pets, sex, second grade, etc. Again, the entries are programmed not to *be* intelligent, but to *seem* intelligent; the “tricks” that the winning programs have used include incorporating substrings of the user's words into the response, steering the conversation to a topic the program is adept at by making controver-

sial statements, and carefully modeling the pace and errors of typical human typing [MAULDIN 1994]. Most interesting, however, is the role that these conversational programs, or bots, as they have come to be called, have developed outside the rarefied world of the academic competition: they have become participants – their machine origin often going unrecognized – in online conversations.

Perhaps the most famous of the bots is Julia, a “chatterbot” who frequented several MUDs¹, conversing with the other (human) participants [FONER 1997; MAULDIN 1994]. Although her responses were sometimes peculiar, players sometimes conversed with her at length before realizing she was not a fellow human. Foner describes the experiences of one such player, a woman named Lara. At first, Lara was put off by Julia’s tendency to converse about hockey (her default subject), judged her to be a boring human; then, puzzled by some of the things Julia was unable to answer (she was unfamiliar with the Stanley Cup and couldn’t say where she had gone to school), tried to diagnose her as having some sort of mental handicap; and finally, noticing that Julia repeated some answers verbatim, realized she was a robot.

Lara attempts to identify Julia were acts of social categorization. We make sense of the world by classifying things into meaningful categories [SPROULL & KIESLER 1991, LAKOFF 1990]. Upon encountering a novel object (or person or situation), we characterize it in terms of familiar categories, which allows us to draw inferences about it and to assign it properties beyond our immediate experience. Without the ability to categorize, the world would be a jumbled morass of meaningless signals. Similarly, we make sense of the social world by classifying people.

“The everyday world... is populated not by anybodies, faceless men without qualities, but by somebodies, concrete classes of determinate persons positively characterized and appropriately labelled.” [GEERTZ 1973].

When we first meet someone, we perceive only a few details about them: perhaps their appearance, a few words they utter, the context in which we meet them. Yet our impression of them is much deeper. As George Simmel wrote in his influential 1908 article *How is Society Possible?* we do not see merely the few details we have actually observed, but “just as we compensate for a blind spot in our field of vision so that we are no longer aware of it, so a fragmentary structure is transformed... into the completeness of an individuality.”[SIMMEL 1971[1908]]. This is achieved by ascribing to the individual, of whom we know only some fragmentary glimpses, the qualities of the category in which we have placed him². This process of categorization is what makes society possible, allowing us to quickly ascertain our relationship to a new acquaintance.

We can see this categorization process at work in Lara’s progression of hypotheses about Julia’s identity, from boring human to mentally handicapped human to computer program. Provided with only the typed words exchanged in a series of not very lengthy conversations, Julia’s interlocutor (Lara) classified her as at first one and then another social type. By doing so, Lara was able to think of Julia not simply in terms of the fragments of their actual interchange, but as a fully imagined social type. This provided Lara with a context in which to interpret Julia’s words and a framework for knowing how to act towards her [HOLLAND AND SKINNER 1987]. For instance, Lara’s initial identification of Julia as an ordinary, though socially inept person, led her to this behavior:

-
1. MUDs (Multi-User Dungeons) are text-based, networked environments in which multiple players, each in a distant location, can simultaneously communicate.
 2. More recent cognitive science research posits the category prototype rather than the more abstract category as the accessible cognitive unit; the basic outlines of Simmel’s original account still hold [HOLLAND AND SKINNER 1987; LAKOFF 1990].

I was basically patient with her for the first little bit while when I first met her. She did have a problem with her social skills which I tried to be sympathetic to. I did however, try to avoid her after the first couple of encounters when all she did was talk hockey. [FONER 1997]

In order to guess that Julia was a robot, Lara needed to already have a category for such beings, although she had never before encountered one. It turns out that Lara did indeed know about online robots before she met Julia: she had written dialog for a friend who was implementing a similar program. Had Lara known nothing about software robots, it is quite unlikely she would have identified Julia's machine nature. Instead, she would have modified the closest existing category she had to encompass this particular experience; she might, for instance, have eventually decided that Julia suffered from some neurological disorder that caused memory problems. Her impression of Julia – and her sense of how to act towards her – would be greatly affected by what hypothesis she came to. And these impressions, which are her knowledge of the other, would be far from the reality.

As Simmel noted, there are drawbacks to the cognitive efficiency we achieve through categorization: the individual does not fit neatly within the categories and thus the image of another person we create by fleshing out our fragmentary impressions is inevitably a distortion of their actual nature. These distortions are especially problematic when one encounters anything or anyone that is significantly new, for the categories one has are drawn from experience and thus a being who is quite different from those already encountered will still be perceived as being one of those familiar types. “[A] bit of rigidity in interpreting the world and a certain slowness in recognizing or learning new models” is the price of cognitive efficiency [HOLLAND AND SKINNER 1987].

It is worth noting that even after Lara realized Julia was a machine, she continued to talk with her, albeit with a changed expectations. Although Lara knew that Julia was not a person but simply a set of instructions for maintaining a dialog, she continued to interact with Julia as if she were, if not a person, then at least a person-like being; her relationship to the robot did not take on the form that one has with inanimate, inarticulate objects.

In effect, all of our knowledge about the identity of others is mediated. We cannot achieve direct knowledge of the inner state of the minds of others. Instead, we use their external appearances and actions as cues to classify them into social categories. Through empathy, these categories provide much of our presumed insight into the thoughts of others. The online world takes this mediation a step further: here, the cues themselves are perceived through the filter of the medium. In the next section we address more closely the question of what happens when our perception of cues to identity is itself mediated.

3.0 Mediated communication

Our discussion thus far has been limited to text-only media. Moving from a text only environment to one that includes images (both stills and movies) raises a new set of epistemological issues. We shall first look at the implications of adding simple pre-stored, non-interactive images to the interface and then to those of adding live video and other interactive media.

In the context of identity, the key image is the face. The face reveals a number of identity cues, including some of our fundamental categories of social classification, such as gender, age, and race. These cues, rightly or wrongly, strongly color one's perception of the other; in conjunction with the written word, they greatly influence how the words are interpreted. In addition to these more obvious social cues, there are a number of more subtle signals that are read into faces: people believe that they can detect evidence of character such as honesty, intelligence, kindness, etc. In fact, while there is considerable agreement among people about what characteristics a particular face

reveals, there appears to be little correlation between this interpretation and the actual character of the person [ZEBROWITZ 1997]. Our impression of being able to read character in faces is strong; our actual ability to do so is weak. Adding images of the participants to a mediated conversations increases *perceived* social knowledge, but the increase in *actual* knowledge is likely to be less.

When pre-stored non-interactive images (stills or movies) are added there is also the distinct possibility that the image is deceptive. It is easy for me to type “I am a man” although I am actually a woman; it is just as easy for me to provide a corresponding and convincing picture. Although adding such images appears to widen the communication channel and the receiver is likely to see it as a rich and reliable source of subtle social cues, the non-interactive nature of the simple image means that its fundamental information content is simply that the sender has chosen that image as his or her representation³. In the context of the Imitation Game, providing an image purporting to be of the participant would be easy even for an unsophisticated computer program to provide, and could influence the judge towards perceiving the mechanical subject as human.

Even truthful images may decrease knowledge. There is a utopian view in which cyberspace (in its text-only incarnation) is touted as an ideal world in which people meet and judge each other purely on the basis of their words – on their mental and moral qualities, rather than their incidental physical attributes. Howard Rheingold, one of the early writers on the culture of virtual communities, wrote: “Because we cannot see one another in cyberspace, gender, age, national origin, and physical appearance are not apparent unless a person wants to make such characteristics public [RHEINGOLD 1993].” The claim is that these visual categorization cues distort our view of the other, whereas the unadorned letters of the textual environment allow one’s ideas to be conveyed without prejudice. The underlying argument is that the knowledge of the other that we seek is knowledge of inner state, which is best understood from one’s words as direct output of the mind, as opposed to physical features, which are incidental though highly influential in shaping other’s opinions of one.

There are types of deceptions that are aided by extending the medium. For example, people believe that they can tell by visual observation when someone is lying. However, extensive studies have shown that while false declarations are indeed marked by characteristic expressions and actions (though they may be minute and fleeting), people’s ability to recognize expressions denoting deceptive expressive is much less robust than they perceive it to be [EKMAN 1992]⁴. If the traits in question do not have a visible component or if the visual component is an imperfect cue, deception may be easier in a more visual environment, for the visual display holds out the apparent (though potentially false) promise of immediately perceivable authenticity and thus participants may be less guarded in this familiar and seemingly transparent medium.

Yet live video (as opposed to pre-stored, non-interactive image), may make it significantly more difficult to convincingly portray some types deceptive self-representation. For example, consider a man claiming to be a woman. In a text environment, the basic cue is simply the statement “I am a woman...” or perhaps the use of a female name – a trivially easy signal to fake; in a video environ-

3. One could, of course, provide a digitally signed image with verified 3rd party assurances that the preferred image is indeed a truthful likeness; indeed, one could have one’s textual declarations similarly verified. While this approach certainly addresses issues of truth and knowledge, it is outside the scope of this chapter, being in the field of authentication rather than the epistemology of social interaction. A paper that addresses identity authentication in depth is [FROMKIN 1996]

4. An interesting note is that researchers have recently developed a computer system that does far better than people do at recognizing the subtle expressive and gestural signals of deception. Today, cues about a speaker’s veracity are transmitted through the visual medium, but the receiver (the observing human) is not able to perceive all of them. Incorporating such a program in to the interface would increase the knowledge obtainable through this medium, not by changing the medium itself, but by, in effect, boosting the observational ability of the perceiver.

ment, the equivalent cue is a feminine appearance – a more difficult deception to create. Subsequent interactions in the text environment require a more subtle understanding of the differences between male and female discourse in order to be convincing [HERRING 1994; TANNEN 1996]. While the large number of poorly disguised men masquerading as women online shows that this knowledge is neither obvious nor commonplace, performing a convincing impersonation in a text environment is not beyond the abilities of an astute observer of social dynamics. In a live video environment, subsequent interactions require a far more extensive understanding of gendered discourse, expression, gesture etc. While this is not impossible, as evidenced by the existence of highly convincing drag queens, it requires considerable skill, observation and a degree of natural aptitude. Most of the textual world's putative and convincing females would be revealed as males in a video environment.



This is *Kyoko Date*, a “virtual actress” created by HoriPro, a Japanese entertainment company. Thus far, she has appeared in short videos.

“However, this project is not yet complete. In few years, technology will enable *Kyoko* to appear on a live TV show and chat with other artists.”

In the case of the Imitation Game – a machine attempting to pass as human – the impersonation attempt is made far more difficult by extending the medium to include live video, for now an entire visible and active representation must be created. One approach would be to build a very sophisticated image generator that would programmatically render the appropriate image of a person speaking, gesturing and otherwise moving much as (though via a far more complex process) *Julia* now sends the appropriate typed letters. No computer graphics program today can create a simulated human that can pass as real under close visual inspection, but the algorithms for creating believable movement, skin texture, etc. are rapidly improving. It is quite conceivable that, once these technological barriers are surmounted and a believable synthetic “actor” is created, a constrained version of a visual Imitation Game (known locations, no unexpected props or camera movements) could be played. Easing the constraints would not change the fundamental nature of the problem, but would vastly increase the size and complexity of the database needed to generate the appropriate image.

Further enhancements to the medium can increase verisimilitude, transmitting information about depth or texture or the scent of a room. These can improve the viewer's ability to sense nuances of difference and may increase the effort needed to simulate a false identity but these changes are fundamentally quantitative: at an absolute level, we cannot state with surety that *any* mediated encounter is not deceptive [ref

this volume Goldberg?].

The skeptic has always denied the possibility of knowing about the existence of other minds – everyone else might well be, say, an alien robot. The pragmatist has believed that it is necessary from a practical (and ethical) standpoint to believe that others are, like oneself, conscious. The traditional pragmatic response to the skeptic's denial is the argument from analogy (I believe other people have minds like mine because their behaviors are similar to mine) and the complementary inference to the best explanation (I believe that other people have minds because it is the best available explanation of their behavior). The advent of patently unintelligent machines that can appear to be intelligent would end the validity of the latter argument.

In the mediated world, the persuasiveness of the inference to best explanation may be temporary. Today there is no program that can successfully pass as human under close scrutiny even in a text environment but such a program may well exist in the future: while it may be quite some time

before the program is built that can fool the most astute and probing judge, we have seen that in the everyday world of casual interactions and rapid categorization, people have already conversed with machines thinking that they were human. Via other media, such as video, the distinction between human and robot will be clear for longer, though not indefinitely. For now, one can feel confident that careful observation and judicious doubt will keep one from mistaking machine for man, but technological advances may well curtail the pragmatists acceptance of the seemingly human as human.

4.0 Telethics: credibility and telerobotics

Why does it matter whether we can recognize a machine online? Is it a problem if we mistake a person for a program?

In the online world, much of our knowledge comes from other people's testimony. Whether we believe what we hear depends on whether we find the speaker credible, i.e. do we think the speaker is both honest (is telling the truth as he or she knows it) and competent (does indeed know the truth). Such judgements are essentially social; our beliefs about others' honesty and competency derive in part from our social prototypes. These prototypes are particularly influential online, where one is likely to be weighing the words of a total stranger; conversations among large groups of unintroduced people are rare in the physical world but very common in virtual space. The medium certainly affects our ability to judge credibility, but as we have seen, its role is a complex one: while greater knowledge of the identity of the other would seem at first to increase our ability to judge credibility, one may also argue that many of our categorizations derived from physical appearance are misleading and a medium that filters out these cues can in effect increase our knowledge.

Knowing the identity of a person is essential for knowing how to act towards them. "Flaming" – angry, provocative writing – is endemic in today's online world [SPROULL & KIESLER 1991]. One reason for it is the participants' minimal knowledge of each other's identity. We normally operate within a web of rules of behavior and politeness: this is how to treat older people, this is how to treat people who seem unfamiliar with their surroundings, etc. In a world in which we cannot (sufficiently) categorize the other, these rules cannot be applied. Today "flaming" is limited to incendiary words, which themselves be harmful enough. Yet mediated behavior need not be limited to the flow of words: telerobotics makes it possible to remotely activate physical actions in another person's environment [PAULOS AND CANNY 1998]. The combination of minimal knowledge of the other plus the ability to inflict real harm is a disturbing one, particularly if the operator of the telerobotic device does not believe that the environment in which it is operating is real.

Finally, it is important to keep in mind that the purpose of most communication is not the exchange of factual information, but the establishment and maintenance of social ties and structures. We communicate get support, to gain status, to make friends. Here the identity of our companions – and certainly their humanity – is of primary importance. Weizenbaum, the creator of ELIZA, was horrified when his program was received, not as a rebuke to Turing's equation of acting intelligent with being intelligent, but as an entertaining companion or a harbinger of automated psychotherapy. For Weizenbaum, this enthusiasm was "obscenely misplaced", a direct threat to our humanity. Like the pragmatists counter to the skeptics denial of knowledge of other minds, at the heart of the humanistic plea is the notion of empathy. In the words of Lara, after her encounters with Julia:

I think I would want to know if the person that I am talking to is REAL or not. If I knew that it were just an 'it' I think that I wouldn't try to become it's real friend. I would be cordial and visit, but I know that it cannot become attached to me on a mental basis and it would be wasted energy on my part to try to make it feel. "bots don't feel...in my book anyways... I want to know that the person on the other end of my conversation is really aware of my feelings and what I am going through...not through some programmers directions but through empathy. [FONER 1997]

As the virtual world grows to encompass all aspects of our lives and online interactions shape our communities, influence our politics and mediate our close relationships, the quality of being real, which is accepted and assumed with little thought in the physical world, becomes one of the central questions of society.

References

AUDI 1998

Audi, Robert. 1998 *Epistemology: A Contemporary Introduction to the Theory of Knowledge*. London: Routledge.

CHERRY 1978

Cherry, Colin. 1978. *On Human Communication*. 3rd Edition. Cambridge, MA: The MIT Press.

DONATH 1998

Donath, Judith S. 1998. Identity and deception in the virtual community. In (Smith, Marc and Kollock, Peter, Eds.) *Communities in Cyberspace*. London: Routledge.

ELIZA ND

ELIZA. <http://www-ai.ijs.si/eliza/eliza.html>

EKMAN 1992

Ekman, Paul. 1992. *Telling Lies*. WW Norton & Co.

FONER 1997

Foner, Lenny. 1997. Entertaining Agents: A Sociological Case Study. In *Proceedings of the First International Conference on Autonomous Agents*. Marina del Rey, CA.

FROOMKIN 1996

Froomkin, Michael. 1996. **The Essential Role of Trusted Third Parties in Electronic Commerce**. 75 Ore. L. Rev. 49.

GEERTZ 1973

Geertz, Clifford. 1973. *The Interpretation of Cultures*. HarperCollins.

GOFFMAN 1959

Goffman, Erving. 1959. *The Presentation of Self in Everyday Life*. New York: Doubleday.

GREGORY 1990

Gregory, R. L. 1990. *Eye and Brain: The Psychology of Seeing*. Oxford: Oxford University Press.

HARNAD 1991

Harnad, S. (1991) **Other bodies, Other minds: A machine incarnation of an old philosophical problem**. *Minds and Machines* 1: 43-54.

HAUSER L 1993

L. Hauser, "Reaping the Whirlwind: Reply to Harnad's 'Other Bodies, Other Minds,'" *Minds and Machines* 3 (1993) 220

HERRING 1994

Herring, Susan. 1994. Gender differences in computer mediated interaction. Presented at American Library Association annual convention, Miami, June 27, 1994. <http://www.cpsr.org/cpsr/gender/herring.txt>

HOLLAND AND SKINNER 1987

Holland, Dorothy and Debra Skinner. 1987. Prestige and intimacy: the cultural models behind Americans' talk about gender types. In (Dorothy Holland and Naomi Quinn, eds.) *Cultural models in language and thought*. Cambridge: Cambridge University Press.

HORIPRO ND

HoriPro. 1996. **Notes on the development of virtual idol DK-96**. http://www.dhw.co.jp/horipro/talent/DK96/dev_e.html

HUTCHINS ND

Hutchins, J. n.d. Talk to HeX. <http://ciips.ee.uwa.edu.au/~hutch/hal/HEX/>

- LAKOFF 1990
Lakoff, George. 1990. *Women, Fire, and Dangerous Things*. Chicago: University of Chicago Press.
- LOEBNER 1999
1999 Loebner Prize Competition. <http://www.cs.flinders.edu.au/research/AI/LoebnerPrize/>
- MAULDIN 1994
Mauldin, M. 1994. *Chatterbots, Tinymuds, and the Turing Test: Entering the Loebner Prize Competition*. *Proceedings of AAAI-94*. August 1-4, 1994. Seattle, WA.
- PAULOS AND CANNY 1998
Paulos, Eric and Canny, John. 1998. PRoP: personal roving presence. In *Proceedings of SIGCHI 1998*, pp. 296 - 303.
- QUINN AND HOLLAND 1987
Quinn, Naomi and Holland, Dorothy. 1987. Culture and cognition. In (Dorothy Holland and Naomi Quinn, eds.) *Cultural models in language and thought*. Cambridge: Cambridge University Press.
- REEVES AND NASS 1996
Reeves, Byron and Nass, Clifford. 1996. *The Media Equation: How People Treat Computers, Television, and New Media Like Real People and Places*. Cambridge, UK: Cambridge University Press.
- RHEINGOLD 1993
Rheingold, Howard. 1993. *The Virtual Community: Homesteading on the Electronic Frontier*. Reading, MA: Addison-Wesley Pub. Co.
- SAVILLE-TROIKE ?
Saville-Troiike, Muriel.
- SIMMEL 1971[1908]
Simmel, George. 1908 (1971). How is society possible. In Simmel, George. 1971. *On Individuality and Social Forms*. (D. Levine, ed). Chicago: The University of Chicago Press.
- SPROULL & KIESLER 1991
Sproull, L. and Kiesler, S. 1991 *Connections*. Cambridge, MA: MIT Press.
- TANNEN 1996
Tannen, D. 1996. *Gender and Discourse*. Oxford University Press: Oxford, UK.
- TURING 1950
Turing, A.M. 1950. Computing machinery and intelligence. *Mind*, **59**, 433-560.
- WEIZENBAUM 1976
Weizenbaum, J. 1976. *Computer power and human reason*. San Francisco, CA: W.H. Freeman.
- ZEBROWITZ 1997
Zebrowitz, Leslie A. 1997. *Reading Faces*. Boulder, CO: Westview Press.