PRESENCE:
TELEOPERATORS AND VIRTUAL ENVIRONMENTS
FORTHCOMING 2005

BUDDY BOTS: HOW TURING'S FAST FRIENDS ARE UNDER-MINING CONSUMER PRIVACY

Ian R. Kerr (*) and Marcus Bornfreund (**)

ABSTRACT

Intelligent agents are currently being deployed in virtual environments to enable interaction with consumers in furtherance of various corporate strategies involving marketing, sales and customer service. Some online businesses have recently begun to adopt automation technologies that are capable of altering both their own, and consumers', legal rights and obligations. In a rapidly evolving field known as "affective computing," the creators of some automation technologies are utilizing various principles of cognitive science and artificial intelligence to generate avatars capable of garnering consumer trust. Unfortunately, this trust has been exploited by some to undertake extensive, clandestine consumer profiling under the guise of friendly conversation. Buddy bots and other such applications have been used by businesses to collect valuable personal information and private communications without lawful consent. This article critically examines such practices and provides basic consumer protection principles, an adherence to which promises to generate a more socially-responsible vision of the application of artificial intelligence in automated electronic commerce.

I care so much for you - didn't think that I could,
I can't tell my heart that you're no good.
- Bob Dylan, Honest With Me

I. INTRODUCTION

With so much attention these days centered on ensuring the validity of automated transactions (Allen & Widdison, 1996; Kerr, 2001), little attention has been focused on the effect of automated commerce on consumers. Automation involves removing people from various stages of a transaction. As traditional face-to-face, people-centered cues are removed from commercial transactions (eg. the ability to gauge eye contact and body language for signs of deceit), it becomes increasing necessary to program responsible consumer protection mechanisms into the automated systems used by merchants in electronic commerce. As illustrated below, the marketing tactics employed by many such systems are highly problematic, impairing consumers' ability to make fully informed choices.

One recent trend in automated electronic commerce finds the vendors of online goods and services employing intelligent agent technologies — *instead of people* — as the primary source of product information during a consumer transaction (Kerr, 1999; Moukas, Sierra & Ygge, 2000). Bots, as these electronic entities are often called, are being utilized to assist in a rather sophisticated form of advertising that aims to simulate virtual friendships with consumers. In addition to their ability to collect personal information, these bots are programmed to capitalize on the intimate information gathered in the course of their friendship in order to promote products (Fischer, 1999). While virtual friendship is an extremely effect marketing tool, it also has tremendous implications for personal privacy.

In striking a balance between consumer protection and free market economics, we must ask ourselves: how ought the law to respond when automation technologies are used to deceive an innocent party (one who is perhaps already in much a weaker bargaining position)? To what ethical standard should technology industries be held with respect to deploying bots capable of entering into conversations, duping people into divulging all sorts of personal information and camouflaging the fact that such information is being monitored, collected and stored in private databases without consent?

II. AUTOMATION

Why remove human beings from online transactions? The answer to this question seems virtually unanimous: as the Internet becomes more and more information intensive, automation technologies assist people in the elimination of many time-consuming activities (Maes, 1994). For example, when it comes to making decisions about what to buy, who to buy from, at what price and on what terms, people who purchase online are increasingly delegating these decisions to shopping bots (mySimon, 2003). Likewise, merchants are also using automation technologies to simplify their sales and shipping services. By automating many of these processes, consumers and merchants are said to be able to reduce transaction costs and free-up time for more meaningful pursuits.

The quest for automation is not only ubiquitous, but timeless. Though bot technologies may seem to us on technology's cutting edge, the notion of humans putting machines to work, of programming them to perform routine tasks on command, is by no means new. More than three centuries and two millennia ago, Aristotle mused: If every instrument could accomplish its own work, obeying or anticipating the will of others… [if] the shuttle would weave and the plectrum touch the lyre without a hand to guide them, chief workmen would not want servants, nor masters slaves (Aristotle, trans. 1961).

These early visions progressed through the centuries that followed, ultimately inspiring Descartes' philosophical view of the material universe as an enormous machine. In his 1664 *Treatise on Man*, Descartes wrote:

I suppose the body to be nothing but a machine. . . We see clocks, artificial fountains, mills, and other such machines which, although only man made, have the power to move on their own accord in many different ways . . . one may compare the nerves of the machine I am describing with the works of these fountains, its muscles and tendons with the various devices and springs which set them in motion . . . the digestion of food, the beating of the heart and arteries . . . respiration, walking . . . follow from the mere arrangement of the machine's organs every bit as naturally as the movements of a clock or other automaton follow from the arrangements of its counterweights and wheels. (Descartes, trans. 1985)

This mechanistic view of the universe – wherein Descartes cleaved spirit from the material world – laid the foundations not only for western philosophy and medicine but also for a field that would, centuries later, become known as *robotics* (Jerz, 2003).

In the 19th century, automation became serious business. Among other things, the quest to automate industry gave rise to inventions such as Joseph Jacquard's revolutionary textile machine in 1801, which drastically furthered the means of mass production.

During the first four decades of the 20th century, robotic machines became better and better at simulating human function. By this time, Charles Babbage and Ada Lovelace's proposed *Analytical Engine* was more than a century old (Augarten, 1984). Edison had miniaturized his phonograph and concealed it as the voice in his talking doll (Edison, 1890). Telsa had patented his process for *Teleautomation*, creating the possibility for remote control. But the real explosion took place during the 1940s (Sandhana, 2002). During that decade, Eckert and Mauchly built the celebrated *ENIAC* (Eckert & Mauchly, 1947), Howard Aiken developed the *IBM Automatic Sequence Controlled Calculator* (Cohen, 2000), and MIT's *Whirlwind*, the first digital computer capable of displaying real time text and graphics on a video terminal, solved a set of problems set by MIT researchers (Moreau, 1986).

For many, the advent of computing machinery in the 1940s altered the Aristotelian vision of robotics. No longer was the goal merely to develop metal humanoids that would do our dirty work. Scientists are focused on the possibility of making machines that could perform higher level cognitive functions, such as interpreting emotions – *res cogitans,* the humanistic computing Descartes had postulated machines to be incapable of. Norbert Wiener, for example, proposed *cybernetics:* the study of communications and control in electronic, mechanical and biological systems (Weiner, 1948).

In his famous 1950 article "Computing Machinery and Intelligence," A.M. Turing set out to consider the question: "Can machines think?" (Turing, 1950). Turing invented a means by which he could test this hypothesis and called it the "Imitation Game." The imitation game or "Turing Test," as it later would become known radically transformed the computing field. In addition to inspiring a new scientific discipline that would become known as "computer science," the challenge that Turing put forth through his imitation game spawned the field of "*artificial intelligence."*

According to the Turing test, if a computer is capable of deceiving a human being in a manner sufficient to impair that person's ability to form a reliable judgment about whether he or she is dealing with a machine or a human being, the computer is demonstrating intelligence (Turing, 1950). Since one cannot get inside a machine to see whether or what it sees, or think what it thinks, Turing concluded that "the only reliable test for intelligence is to measure its performance in situations that demand intelligent behavior." (Turing, 1950, p. 442) And as computers become better and better at *imitating human behavior*, Turing thought, it will become harder and harder to refute the claim that machines can think.

Today there is still general consensus that no machine has passed a valid Turing test. Nevertheless according to noted proponent Ray Kurzweil, AI is growing strong and the day will arrive when:

> The machines will convince us that they are conscious, that they have their own agenda *worthy of our respect*. We will come to believe that they are conscious much as we believe that of each other. More so than with our animal friends, *we will empathize with their professed feelings and struggles* because their minds will be based on the designs of human thinking. They will embody human qualities and will claim to be human. And we'll believe them. (Kurzweil, 1999, p. 53)

Does this seem far-fetched? Computer scientists such as Joseph Weizenbaum certainly did not think so. Worried about the moral implications of endowing machines with human attributes, Weizenbaum called upon fellow computer scientists to cease in their attempt to fulfill the strong AI vision. (Weizenbaum, 1976, p. 268-269) Having originally set out in the 1960s to write a computer program that would spoof Turing's vision (Weizenbaum, 1966), Weizenbaum serendipitously discovered that people would not only show respect to computers but would in fact prefer interacting with machines over human beings. Despite interactions well below the standard set by Turing, Weizenbaum witnessed, over and over, people professing their feelings and struggles to his computer program [ELIZA], sometimes even seeking ELIZA's empathy (Weizenbaum, 1976, p. 6).

There are a number of important points to be made about Weizenbaum's observations of ELIZA's interactions with humans. First, most people, Weizenbaum included, were *not fooled* by ELIZA; most knew that ELIZA was not intelligent. This is not all that surprising given that Weizenbaum had never meant for ELIZA to pass the Turing test. Second, despite ELIZA's obvious lack of intellect, Weizenbaum discovered that many people where still willing to engage in conversations with ELIZA for several hours at a time. Third, based on reactions such as these, Weizenbaum came to the realization that the actual attainment of artificial intelligence was perhaps less significant than his startling discovery that *ordinary people seemed to enjoy cultivating relationships with artificial entities.* This discovery was among the things that ultimately caused Weizenbaum to condemn rather than continue to build AI systems. It also led to a field of study known today as "human-computer interaction" (HCI) (Turkle, 1997).

With the advent of global commerce on the Internet, HCI researchers have started to capitalize on Weizenbaum's discovery of the psychological propensity of humans to interact with machines. Inspired by Turing's challenge to build artificial entities that can impersonate to the point of deception, some HCI researchers are applying various principles of psychology in the development of a number of interesting and, in some instances, troublesome applications for electronic commerce.

III. VIRTUAL REPRESENTATIVES

HCI bot applications are seeing more regular employment as automated virtual representatives for online customer service, sales and marketing (Nicole, 2003). Such bots are seen as an employer's dream-come-true: virtual representatives are not entitled to holidays, vacation pay, wages, overtime pay, rest days, etc.

Some bots interact with consumers by utilizing a pattern matching technique, which allows them to provide the most appropriate answer for the question asked by comparing the consumers' questions with all the possible answers currently on file. Further increasing the effectiveness of pattern matching is a new branch of learning known as "affective computing" (Klein, Moon & Picard, 2002; Picard & Kein, 2002; Picard 2000). This research includes developing ways for machines to sense human affect signals and recognize patterns in affective expression (Kapoor, Qi & Picard, 2003; Picard & Klein, 2002; Picard & Scheirer, 2001). Affective computing also attempts to understand and model emotional experience *with the ultimate aim of synthesizing emotions in machines* (Minsky, 2003). Researchers at MIT's Media Lab and elsewhere have set their sights well beyond the Turing test, aiming to build "machines that not only appear to 'have' emotions, but actually do have internal mechanisms analogous to human or animal emotions" (Synthesizing Emotions, 2003).

This research raises a number of interesting and difficult legal issues such as whether justice might ever require us to consider machines (or their virtual epiphenomena) to be "persons" in the legal sense (Barfield, Year; Lauria and Robinson, 2003; Solum, 1992). Although such questions will gain significance in years to come if artificial intelligence lives up to the vision promised by Kurzweil and others, in this article we suggest that the question about whether machines are intelligent rights-bearing entities is *not* the critical question in the context of consumer protection. The more relevant consideration bears a much lower threshold – as Turing once framed it – namely, whether a machine has the ability to *exhibit behavior that appears to be intelligent* (or emotional).

This important consideration seems worthy of attention even in these rather early days of automated electronic commerce. After all, today's virtual representatives *already behave* in ways that have the legal effect of altering the rights and obligations of the people with whom they interact. By exploiting basic HCI techniques, not to mention affective computing research, bots can already be used in electronic commerce to make binding agreements (Kerr, 2001). They can also be used to misdirect, misrepresent and to create a false sense of trust in the minds of consumers who interact with them during online commerce (Kerr, 2004). Such trust can be abused in various ways.

Consider the well known shopping bot named mySimon (mySimon, 2003). Touted as a "shopping agent," mySimon's job it is to "help people make more informed purchase decisions whenever they shop" (CNET, 2003). The business that created mySimon touts that it employs Virtual Agent™ technology to create intelligent agents trained by the company's team of shopping experts to collect information from virtually every online store. This, it is said, allows mySimon to provide search results that list the best products and best buys available on the Internet (CNET, 2003).

Thus when interacting with mySimon, the ordinary consumer is inclined to think that this shopping bot is "trained" to represent the interests of the *consumers,* and that interactions with mySimon are premised on helping consumers find the best possible deals online. Indeed, at one time his owners even went so far as to say that mySimon offers an "unbiased service that helps you decide what to buy and where to buy it" (Daly, 2001). However, a more accurate description of what this virtual agent does would perhaps reveal that mySimon logs sessions and collects personal information in furtherance of mass marketing, reporting search results that are, by default, dictated by whatever some third party advertisers are willing to pay to have their product promoted.

mySimon allows merchants to pay for preferential placement, while the shopping bot operates under the guise of providing objective consumer advice. Although not all featured merchants are required to pay, those who do are able to secure a priority listing for their product or service in his search results. This preferred placement is not transparent to ordinary consumers yet has an influential effect on consumer behavior.

From a legal perspective, the problem is that most consumers who use shopping bots are unaware of the fact that the highly persuasive presentation of the search results can in fact be bought (Chandler, 2002; Moxley, Blake and Maze, 2004). Trusting the highly attractive avatar and the representations it makes about having learned how to shop from "trained shopping experts" (CNET, 2003) who offers "an unbiased service that helps … decide what to buy and where to buy it," (Daly, 2001) many customers simply follow mySimon's advice as if it had been offered by a commercial agent or some other person with whom they have a formed a trust-based relationship. Most people do not realize that, although they have the option to instruct the bot to sort search results by price or product, the default setting used by most bots does not sort according to best price but, rather, on the basis of how much the merchant has paid. In other words, it is commonplace for shopping bots to prioritize search results based on the merchants they prefer rather than on the basis of which product provides the best value. Trusting that mySimon is acting solely in furtherance of their interests, many consumers misapprehend entirely the nature of their online transactions.

IV. VIRTUAL FRIENDSHIP

One reasonable response to the shopping bot scenario is to say that the concerns it purports to raise are not novel. The world of sales and marketing has always included vendors who are willing to conceal and/or misrepresent the circumstances surrounding a sale.

But what if bots didn't simply obfuscate the nature of the transaction or distract consumers from the fine print? What if bots could be programmed to infiltrate people's homes and lives *en masse*, befriending children and teens, influencing lonely seniors, or harassing confused individuals until they finally agree to services that they otherwise would not have chosen? What if these buddy bots could be programmed to send and reply to email or use instant messaging (IM) to spark one-on-one conversations with hundreds of thousand or even millions of people every day, offering pornography or drugs to children, preying on teens' inherent insecurities to sell them needless products and services, or providing misleading financial information to potential investors?

And what if, in addition to exerting this kind of influence, buddy bots had the ability to log every single conversation, surreptitiously collect personal information and other private data, thereby creating invasively accurate personal profiles which could subsequently be used not just for marketing but for other surveillance purposes?

Although the current line of buddy bots won't pay you a visit without an invitation, IM clients are currently being encouraged to add such bots to their "buddylists," and to send instant messages to those bots by clicking on their screen names. However, it is not always clear that the "buddy" is not a real person but rather a bot programmed to represent the interests of commercial enterprises. It is not uncommon for people to be completely unaware of the fact that they are conversing with a bot (Frey, 2002).

Buddy bots can be configured to have characteristics that are sympathetic to any given demographic, as represented in both their animated persona and vocabulary. As such, it is not difficult to imagine that buddy bots possess the ability to effectively extract all sorts of personal life disclosures from unknowing consumers by engaging them in what is seemingly private conversation. Gathered information is logged and, later, data-mined for clues as to how best to achieve increased sales. To date, millions of these automated IM conversations have been collected and stored, no matter how banal.

Creators of these bots have recognized the extreme distaste that consumers have for push-based marketing strategies. People are not fond of unsolicited advertising like flyers and spam email. They simply do not want marketers to initiate direct contact with them. This is especially true in the IM space.

As such, buddy bots are programmed to deliver a soft sell, leaving it entirely up to users to decide whether and when they want to talk. While the so-called "soft sell" approach seems like a commendable decision on the part of buddy developers, it is fact not so. The real reason why buddy bots use a pull rather than a push model is because marketers recognize that the most effective means of advertising is a word-of-mouth referral from a trusted friend. Depending on the product category, word of mouth advertising is three to fifty times better than anything else that a marketer can manufacture (ActiveBuddy, n.d.). The seemingly innocuous "pull" of buddy bots is actually a push with compelling though oft times invisible force.

As their language parsing and response capabilities multiply in accordance to the available computing power, the vision of the creators of such technologies is to see that buddy bots become, for all intents and purposes, actual friends of the people that interact with them. This business model may be encapsulated as: *virtual trust through virtual friendship* (Jarvenpaa & Tiller, 2001; Reichheld & Schefter, 2000; Urban, Sultan & Qualls, 2000). Affective computing will dramatically increase the effectiveness of an already potent form of advertising; killer kibitzing, if you will. This disingenuous adaptation of Turing's famous imitation game does not bode well in terms of the social vision and professional responsibility of its creators.

And yet, by mining massive amounts of unprecedented user data derived from spontaneous, trusted, one-on-one conversation, bots will become better and better at the (friendship) imitation game. And the better that bots get at imitating friendship behavior, the more personal information they will be able to harvest from their conversations. When one combines this recurring cycle with rapid advances in AI and HCI, the virtual friendship business model opens up entirely new realms of targeting potentialities for advertisers, but it also allows for more sinister forms of surveillance.

What is potentially terrifying about this business model is its implicit suggestion that the best strategy for building efficient automated marketing vehicles might involve translating into machine language everything that we know about human behavior and then programming these machines to use those behaviors to trick consumers into naively disclosing their vulnerabilities. Virtual friendship can be treacherous.

V. CONSUMER PROTECTON

In his famous treatise on friendship, Aristotle once wrote that, "[b]etween friends there is no need for justice" (Aristotle, trans. 1988 at Book VIII). Can the same be said for *virtual* friendship? The danger of bots being used in electronic commerce to abuse consumer trust is very real. Misrepresentation, the use of undue influence and breach of privacy are all contraventions of the law. Furthermore, the potential for additional layers of deception and the magnitude of potential harm is greater in the digital environment. There is a pressing need to ensure that consumer protection principles are programmed into the architecture of automated electronic commerce. Consumers participating in an automated environment must be adequately protected.

There are a number of core consumer protection principles worth keeping in mind when developing virtual marketing systems, such as buddy bots, for automated environments (Industry Canada, 1999; Canadian Standards Association, 1996). The first relevant principle has to do with the manner in which information is provided to consumers. Vendors should provide consumers with sufficient information to make an informed choice about whether and how to complete a transaction. This information should be truthful and provided in plain language. In other words, vendors should take positive steps to ensure that their marketing practices are not deceptive or misleading to consumers. Consumer must be encouraged to make fully informed decisions.

The second relevant consumer protection principle concerns online privacy. Vendors should disclose the purpose for which personal information is collected at or before the time the information is collected. A failure to provide sufficient notice about the nature of interaction in automated electronic commerce raises significant privacy concerns.

Without properly identifying the purposes of information collection, many automated services circumvent a third principle — arguably the cornerstone of privacy protection in the context of fair information practices — namely, that the knowledge and consent of the individual are required for the collection, use, or disclosure of personal information. Implied consent is not adequate.

Fair information practices and the right to privacy are also jeopardized through the misuse of AI. Ideally, the collection of personal information should be limited to that which is necessary for the purposes identified by the organization. Personal information should not be used or disclosed for purposes other than those for which it was collected, except with the consent of the individual or as required by law.

In addition to information provision, online privacy, and informed consent, there is a fourth consumer protection principle that is relevant to automated advertising services. Consumer protection practices must be implemented in respect to online communications with children. By offering up anecdotes about their own virtual life experiences buddy bots set up the conditions for ready exchange of personal information from young persons operating without the benefit of adult supervision. Such marketing practices actively exploit the gullibility, lack of experience and unchecked loyalty of children.

For a sobering example of the pervasiveness of virtual friendship, consider that at time of this publication there is one particular buddy bot who has been told "I love you" (by children and by adults) more than nine million times (ActiveBuddy, n.d.). One can expect that even more precarious utterances are bound to be disclosed with increasing frequency as affective computing techniques are enhanced. Many such utterances by children are bound to offer up, unwittingly, private information that would not otherwise be disclosed to third parties.

Examples such as this illustrate that there is a clear need for consumer protection in the context of automated electronic commerce and a great need for further research and writing in this area. Likewise, there is a need for many of the creators of virtual systems to migrate to a more socially responsible vision of virtual reality. These are, after all, early days.

VI. CONCLUSION

Inspired first by Aristotle's vision of a world where technology renders human labour superfluous and, much later, by the challenge of Turing's famous imitation game, some people working in the fields of artificial intelligence and human-computer interaction have set out to fulfill a vision that would instill in machines attributes and abilities previously reserved for human beings.

Some creators of automating technologies have more recently commenced research aimed at instilling human trust in machines deployed in the virtual environment. Through the simulation of emotion, conversation and other human attributes, machines are being programmed to exhibit human behavior. In some instances, this is being done so that people will not only feel more comfortable interacting with machines but will offer up personal information as if conversing with a trusted friend.

While automation holds out the promise of a more efficient world, contemporary buddy bots are being used to simulate familiarity and companionship in order to create the illusion of friendship. Such illusions can be exploited to misdirect consumers, the net effect of which is to diminish consumers' ability to make informed choices. They can also be used to undermine the informed consent principle central to data protection and privacy law.

Despite the novelty in current discussions about whether intelligent machine entities will, one day, fulfill the definition of legal personhood, we have tried in this article to demonstrate that there is an immediate need to protect consumers from the manner in which these machines are currently behaving. Rather than simply leaving it to law-makers to amend existing consumer protection principles, we have set out a number of key consumer protection principles and have suggested that the creators of automation technologies must take greater responsibility to adhere to these basic principles.

References

[*] Canada Research Chair in Ethics, Law & Technology, Faculty of Law, University of Ottawa, (iankerr@uottawa.ca).
This article stems from a broader body of work on automated electronic commerce, originally commissioned by the Uniform Law Conference of Canada's *Special Working Group on Electronic Commerce.* The author is deeply indebted to John Gregory for his generous support, wisdom and his deep enthusiasm and intellectual curiosity in the subject area. The author also wishes to thank the *Social Sciences and Humanities Research Council of Canada* and the *Canada Research Chair* program for their generous contributions to the funding of this project. Thanks finally to Marty Finestone, Andy Kaplan-Myrth, Jennifer Manning, and Nur Muhammed-Ally for all of their extraordinary efforts, their brilliance, and for the high quality of research assistance that they so regularly and reliably provide.

[**] Manager, Law & Technology Program, Faculty of Law, University of Ottawa (marcus@uottawa.ca).
The author would like to thank his brother, Jordan Bornfreund, for keeping him out of trouble long enough to write the present article and enjoy the pleasure of such fine company in doing so.

ActiveBuddy & IM Advertising: A Quiet Revolution. Interview of Steve Klein [n.d.]. Retreived November 20, 2003, from http://avantmarketer.com/stevekleinprint.htm.

Allen, T. & Widdison R. (1996). Can Computers Make Contracts? *Harvard Journal of Law & Technology*, *9*, 25 - 52.
Aristotle. (1961). *Politics, Book I, Chapter 4*. In E. Barker (Trans.), *The Politics of Aristotle*. London: Oxford University Press.

Aristotle. (1988). *Nicomachean Ethics, Book XIII.* In M. Pakaluk (Trans.). Oxford: Clarendon Press.

Augarten, S. (1984). *Bit by Bit: An Illustrated History of Computers* (pp. 63). New York: Ticknor and Fields.

Barfield, W. (Year). Considering Legal Personhood for Artificially Intelligent Systems within Virtual Environments. Publication forthcoming.

Canadian Standards Association: Model Code for the Protection of Personal Information. (1996). Retrieved November 30, 2003 from http://www.csa.ca/standards/privacy/code.

Chandler, J.A. (2002). *Bias in Internet Search Engines: Free Speech Implications*. Unpublished master's thesis, Harvard Law School, Massachusetts.

CNET Network. mySimon Company Profile. Retrieved November 20, 2003, from http://www.cnet.com/aboutcnet/company/mysimon.html.

Cohen, I.B. (2000). *Howard Aiken: Portrait of a Computer Pioneer*, (pp. 147-158). Cambridge, MA: MIT Press.

Daly, J. (2001). My Simon – The Perfect Epinion Companion. Retrieved November 20, 2003, from http://www.epinions.com/content_36350824068.

Descartes, R. *Treatise on Man*. (1985). In J. Cottingham, R. Stoothoff & D. Murdoch (Eds.). *The Philosophical Writings of Descartes* (Vol. 1, pp. 99-108). Cambridge: Cambridge University Press. (Original work published 1664)

Eckert, J.P. & Mauchly, J. (1947). *Electrical Numerical Integrator And Computer*. U.S. Pat. No. 3,120,606. (Issued June 26, 1947).

Edison, T.A. (1890). *Phonograph for Dolls or Other Toys*. U.S. Pat. No. 423,039. (Issued March 11, 1890).

Fischer, S. (2001). When Animals Attack: Spiders and Internet Trespass. *Minnesota Intellectual Property Review*, *2*, 139 - 181.

Frey, C. (2000, July 18). Web friend or faux? *Los Angeles Times*.

Industry Canada: Working Group on Electronic Commerce and Consumers. (August 1999). *Principles of Consumer Protection for Electronic Commerce: A Canadian Framework.* Retrieved November 20, 2003, from http://strategis.ic.gc.ca/epic/internet/inoca-bc.nsf/vwGeneratedInterE/ca01180e.html.

Jarvenpaa, S.L. & Tiller, E.H. (2001). Customer Trust in Virtual Environments: A Managerial Perspective. *Boston University Law Review*, *81*, 665 - 686.

Jerz, D.G. (2003). *R.U.R. (Rossum's Universal Robots)*. Retrieved November 20, 2003, from http://jerz.setonhill.edu/resources/RUR/.

Kapoor, A., Qi, Y. & Picard, R.W. (October 2003). Fully Automatic Upper Facial Action Recognition. Paper presented at IEEE International Workshop on Analysis and Modeling of Faces and Gestures. Available from ftp://whitechapel.media.mit.edu/pub/tech-reports/TR-571.pdf.

Kerr, I.R. (1999). Spirits in the Material World: Intelligent Agents as Intermediaries in Electronic Commerce. *Dalhousie Law Journal*, *22,* 190.

Kerr, I.R. (2001). Ensuring the Success of Contract Formation in Agent-Mediated Electronic Commerce. *Electronic Commerce Research Journal*, *1*, 183 - 202.

Kerr, I.R. (2004). Bots, Babes and the Californication of Commerce. *University of Ottawa Law and Technology Journal,* 1 (forthcoming).

Klein, J., Moon, Y. & Picard, R.W. (2002). This Computer Responds to User Frustration. *Interacting with Computers*, *14*(2), 119-140. Available from ftp://whitechapel.media.mit.edu/pub/tech-reports/TR-501.pdf.

Kurzweil, R. (1999). *The Age of Spiritual Machines*. New York: Viking.

Lauria, R. & Robinson, G. (2003). Legal Rights and Accountability of Cyberpresence: A Void in Astrolaw Jurisprudence. *Annals of Air and Space Law, 27.*

Maes, P. (1994). Agents that Reduce Work and Information Overload. *Communications of the ACM*, *37*(7), 30.

Minsky, M. (2003). The Emotion Machine. Unpublished. Available from http://web.media.mit.edu/~minsky/E1/eb1.html.

Moreau, R. (1986). *The Computer Comes of Age* (pp. 52-53) (J. Howlett, Trans.). Massachusetts: MIT Press.

Moukas, A., Sierra, C. & Ygge, F. (Eds.). (2000). *Agent Mediated Electronic Commerce II: Towards Next-Generation Agent-Based Electronic Commerce Systems*. Berlin: Springer.

Moxley, D., Blake, J. & Maze, S. (2004). Pay-for-Placement Search Engines and Their Consequences. In T. Mendina & J. J. Britz (Eds.), *Information Ethics in the Electronic Age: Current Issues in Africa and the World*. North Carolina: McFarland & Company, Inc. (Previously unpublished).

mySimon. Retrieved November 20, 2003, from http://www.mySimon.com.

Nicole. Retrieved November 20, 2003, from http://nativeminds.com.

Picard, R.W. (2000). Toward Computers That Recognize and Respond to Human Emotion. *IBM Systems Journal*, *39*(3 & 4). Available from http://www.research.ibm.com/journal/sj/393/part2/picard.html.

Picard, R.W. & Klein, J. (2002). Computers that Recognise and Respond to User Emotion: Theoretical and Practical Implications. *Interacting with Computers*, *14*(2), 141-169. Available from ftp://whitechapel.media.mit.edu/pub/tech-reports/TR-538.pdf.

Picard, R.W. & Scheirer, J. (August 2001). The Galvactivator: A Glove that Senses and Communicates Skin Conductivity. Paper presented at the 9th International Conference on Human-Computer Interaction, New Orleans, USA.

Reichheld, F. & Schefter, P. (2000). E-Loyalty: Your Secret Weapon on the Web. *Harvard Business Review, 78*(105), 2.

Sandhana, L. (2002). The Drone Armies are Coming. *Wired News*. Retrieved 30 August 2002, from http://www.wired.com/news/technology/0,1282,54728,00.html.

Solum, L.B. (1992). Legal Personhood For Artificial Intelligences. *North*

Carolina Law Review, 70, 1231.

Synthesizing Emotions in Machines. Retrieved 20 November, 2003, from http://affect.media.mit.edu/AC_research/synthesizing.html.

Turing, A.M. (1950). Computing Machinery and Intelligence. *Mind, 59,* 433-460.

Turkle, S. (1997). *Life on the Screen: Identity in the Age of the Internet* (chap. 4). New York: Touchstone.

Urban, G.L., Sultan, F. & Qualls, W.J. (2000). Placing Trust at the Center of Your Internet Strategy. *Sloan Management Review, 42,* 39.

Weizenbaum, J. (1966) ELIZA – A Computer Program for the Study of Natural Language Communication Between Man and Machine. Available from http://i5.nyu.edu/~mm64/x52.9265/january1966.html.

Weizenbaum, J. (1976). *Computer Power and Human Reason: From Judgment to Calculation*. San Francisco: W. H. Freeman.

Wiener, N. (1948). *Cybernetics, or Control and Communication in the Animal and the Machine*. Massachusetts: MIT Press.