# Are Agents an Answer or a Question?

Joseph A. Goguen
Department of Computer Science & Engineering
University of California at San Diego

**Abstract**

This paper first places recent debates about the efficacy and morality of "intelligent agents" into a historical context that includes work of Kant, Babbage, Minsky, and others, and notes the cyclic nature of such debates. It then adds some social and ethical refinements to the debate, and after reviewing certain social science theories, discusses the new concept of narrative-driven navigation and its application to interface design for a distributed theorem proving system. Among our conclusions is that many difficulties with artificial agents arise from a poor understanding of social aspects of human agents.

## 1 Introduction

Agents, or at least proposals for agents, are becoming increasingly common; e.g., see [36]. Although opponents are relatively rare, Ben Shneiderman, Jaron Lanier and others have attacked the idea of "intelligent agents" on various grounds, often of a social nature. Shneiderman has been mainly concerned with the use of agents in user interface design, whereas Lanier has raised larger questions. Some of the targets of their criticism, particularly Pattie Maes of the MIT Media Lab, have responded, meeting certain major criticisms, and also raising doubts about alternative technologies advocated by Shneiderman. Here we review aspects of this debate, place it in a historical context, and enrich it with some social and ethical questions, which we believe deserve more attention. In particular, we argue that many difficulties with artificial agents arise from our present poor understanding of social aspects of human agents, and that sophisticated applications require a blend of interaction metaphors, including agent delegation, direct manipulation, visualization, and navigation. We also sketch our own recent ideas on an alternative approach.

## 2 A Historical Perspective

Throughout the history of computing, many movements have called for more human-like systems and have used anthropomorphic terminology to generate understanding and support. Such movements have often made excessive claims, perhaps misled by their own rhetoric or their (sometimes impressive) partial successes; excessive claims usually lead to unrealistic expectations, followed by disillusion and disappointment, which is then surprisingly often followed by a new version of the same movement, with nearly the same goals but with somewhat improved terminology and technology.

The most obvious example is probably the field of "artificial intelligence" ("AI" for short), which produced many claims and predictions that turned out to be false [5]; the reductionist "logicist" paradigm followed by early researchers like Minsky and McCarthy stressed the logical representation of mental states, attempted (with some success) to colonize adjacent fields such as cognitive psychology and neuroscience, and reaped huge funding (much of it based on projected military applications) and wide publicity (often based on small "demos"); it then crashed in flames, precipitating the so-called "AI Winter," when numerous AI companies and projects were terminated for failure to scale up their demos to genuine applications. However, there were many modest successes scattered among the more spectacular failures, and there was also a rebirth, based on new technologies, such as neural nets and machine learning, which again appeal to analogies with humans.

Among other fields that have followed a similar trajectory are machine translation, the perceptron approach to pattern recognition, web TV, compiler generators, and logic programming. In each case, as well as in classical AI, failure can be attributed in large part to an inadequate understanding of cognitive and social factors involved. For example, early machine translation work assumed that a good lexicon and a good grammar would suffice, whereas researchers are now well aware of the crucial impor-

tance of both context and background knowledge.

The silver lining on these dark clouds is that many of these failures have led to a deeper understanding of the limitations of the technologies involved, and even, though to a regrettably more limited extent, an enhanced humility, and a better appreciation of what it means to be human. (Some of the best work along this line has drawn inspiration from the late Heideggerian combination of pessimism about technology with deep insight into the nature of being human [17].)

An early instance of the cyclic pattern that we are highlighting can be found in the philosophical work of Charles Babbage, who is often considered a founding figure in computer science. In *The Ninth Bridgewater Treatise* [2], Babbage praised the virtues of an early form of industrialization involving heavy regimentation of workers, claiming it to be one facet of an ideal society, and a reflection of the will of God. A similar approach resurfaced much later in "Taylorism" with its imposition of assembly line techniques in many areas of modern work. Such excessive specialization and trivialization of work is today widely considered counter-productive in manufacturing.

## 3 Critiques of Intelligent Agents

Although there are many advocates of intelligent agents, representing a variety of perspectives and applications, with work appearing in many specialized conferences and journals, there are relatively few critics. Perhaps the two most prominent of these are Ben Shneiderman and Jaron Lanier.

Shneiderman, a leader in user interface design, has sharply criticized the use of agents in user interfaces [32, 33], for reasons that include lack of predictability, disenfranchisement of the user, and unclarity of moral attribution. To expand on these points a bit, Shneiderman claims that agents, especially adaptive agents, are by their nature difficult for users to understand and predict, whereas what users want is a sense of control over their software, so that they know it will do what they want it to do; he also claims that if an agent does something the user did not intend, such as destroying a crucial file, or unleashing a virus on the internet, it is unclear who should be held accountable, the user, the manufacturer of the agent, the installer of the agent, the creator of the last patch, or .....

Lanier [25] claims that "the idea of 'intelligent agents' is both wrong and evil." The idea is wrong, he says, because it will not solve the problems that it is supposed to solve, and it is evil because it will "make people redefine themselves into lesser beings." Regarding the first point, one application of agents that has been much advertised is to help users find what they want on the web. But Lanier argues that agents will inevitably trivialize the interests of users, and "deliver an overdose of kitsch," due to taking the same "lowest-common-denominator approach to content that plagues TV." Unfortunately, his view is supported by even a superficial analysis of many commercial web search engines, though there are of course exceptions [11].

Lanier's second point is his main interest. In brief, he argues that people will come to think of agents as being human-like and actually intelligent, and will therefore suffer a reduction of their own autonomy, from having to force their behaviors into the restricted patterns that are allowed by and effective with the agent; as a result they may also come to think of themselves as being like computers.

To add further perspective to this discussion, we may note that the notion of agent used here has origins in the work of the Enlightenment philosopher Immanuel Kant, who assumed the rationality and autonomy of agents, as part of his project to construct a philosophical foundation for Western morality[1] that did not rely upon theology; without realizing it, he also presupposed a Cartesian mind-body dichotomy [20]. Each of these assumptions has been much disputed by modern philosophers (e.g. [19]) and psychologists, but of course this does not mean that they cannot be useful properties of artificial agents. The adjective "intelligent" has a connotation of Kantian rationality, but seems to be slowly going out of fashion, presumably from being used too much and too vaguely, as well as from its association with discredited logicist tendencies within AI. Someday, the word "agent" may suffer a similar fate.

## 4 The Advocates Respond

Perhaps the most basic argument given by agent advocates in response to their critics is that, due to the ever increasing complexity of both systems and the

[1] The situation is quite different for many Eastern moral system, such as Mahayana Buddhism; e.g., see [29, 24].

environments in which they must work, it is necessary for users to abandon at least some control to automatic systems, whether or not one chooses to call such systems agents; in short, agents are inevitable. An often cited example is the use of filtering agents to stem the flood of unwanted email that most of us experience. Shneiderman concedes the necessity of computational assistance, but insists that users must understand and maintain control over any software to which they delegate tasks [34]. He does not oppose automation, but rather he opposes unpredictability, incomprehension, and disenfranchisement.

One of the most prominent defenders of agent technology and ideology is Pattie Maes, head of the Software Agent group in the MIT Media Lab. In [34], which records a debate between Maes and Shneiderman in 1997, Maes accepts Shneiderman's criteria (of predictability, comprehension and control), and says that these are already part of her research agenda. But Shneiderman and Lanier both say they don't know any agents that perform significant tasks and also adequately satisfy the criteria, and they cite prominent failures, such as Microsoft's infamous "BOB," and its recently demoted "Clippy."

Another branch of the discussion concerns the technologies that Shneiderman suggests be used instead of agents, primarily direct manipulation and information visualization, the first of which he has long advocated. Two somewhat inconsistent directions are taken here, to criticize these technologies, and to incorporate them. Regarding the first move, Maes and others have claimed that many domains are too complex for either direct manipulation or visualization, e.g., the world wide web [34]. Regarding the second move, Maes and others say that they want to exploit the best user interface technology they can get, certainly including both direct manipulation and information visualization.

Regarding responsibility for the actions of an agent, in [34], Maes says unequivocally that the user should be held accountable. But in view of examples given above (such as a badly designed agent unleashing a virus), this does not seem correct. It seems to me that in [34], Maes and Shneiderman were following pre-determined agendas and rhetorical strategies, and hence to a great extent were talking past each other, which is of course far from unusual in public debates. This impression is reinforced by [38], which appears to be Maes' most well considered response to the challenges of Shneiderman and Lanier. This article gives interface design guidelines for agent systems, making the point that agents need to obtain the *trust* of users if they are to be successful, and suggesting in some detail how that might be done, unfortunately, without ever confronting the issue of whether or not an agent actually deserves that trust. It would of course be seriously unethical to encourage trust in an agent that does not do what users really want. On the other hand, [38] does admit that a number of important issues concerning agents remain open, including some of an ethical nature.

It's not so clear where all this leaves us, but one tentative conclusion is that we are on the cusp of another twist in the cycle, where agent research is reinventing itself with nearly the same goals (the main change being a bit more emphasis on user satisfaction), and with an enhanced technology that embraces insights from the user interface community. This conclusion is reinforced by the change in terminology advocated by Maes and some other researchers, which is to drop the word "intelligent" [34]. If all this is true, then what remains appears to be mainly an issue of terminology: should we use the word "agent" for the resulting technology? Or more precisely, where should we draw the line, such that only software with sufficiently great autonomy is called an agent? Another conclusion is that ethical issues are not being adequately addressed, especially by the advocates of agent technology.

## 5 Socially Intelligent Agents

I believe the most important issues concerning agents do not boil down to anything as simple as terminology. Raising the stakes by calling for "socially intelligent agents" [28] makes this clearer by moving the debate into deeper questions about human nature and human society, questions that are obscured when the primary focus is on technology, as in the debate between Shneiderman and Maes [34] that was discussed above. Here are examples of some deeper questions that might be asked:

1. How human can agents appear? How human *should* they appear? What are the ethical issues with anthropomorphism?
2. Can agents have emotions? Should they?
3. What effects can an agent's (simulated?) emotions have on users? How can this be exploited,

e.g., in advertising? Is this ethical?

4. Is it necessarily the case that adaptive agents are less predictable? Is predictability always good?
5. Do socially intelligent agents need good models of their users? What needs to be modeled? How detailed do the models need to be? Are models of the user's model of the agent needed? How far should this recursion on models go?
6. Should models of individual users be public? Or should some information be private? If so, how can this be enforced on remote agents?
7. Can agents produce the appearance of qualities like individuality, intelligence, and confidence? Should they? If so, when?
8. What are the inherent limitations of agent technology? What are its greatest strengths?

These are but a few of the possible interesting questions. Most of them can help us to think more deeply about the nature of artificial agents, as well as what it means to be human. In general, the answers are less than fully clear, even to questions that might at first seem obvious; moreover, most of these questions can lead into even deeper questions.

For example, recent research on the physiological basis of memory shows that it is connected to the limbic system, and thus to the emotions, and indeed, to the whole body. So agents without emotions or bodies aren't going to have the same kinds of associative capabilities for their memories that we do (or at least, they could only have them in more awkward ways). But what would it mean to give bodies to software agents? What could an artificial limbic system be like?

Notice that most of our discussion so far has had at least as much to do with the cognitive capabilities of agents as with their social capabilities; we still haven't said enough about what it might mean to be socially intelligent. Here the questions get even more difficult, because this territory is as yet very little explored. We know little about the technical, social and ethical issues that can arise for communities containing both artificial agents and humans. Here are some sample questions:

1. What would it mean for a software agent to be part of a community? Is there a good operational definition?
2. Should agents lie to other agents? If so, when?
3. Should agents lie to humans? If so, when?
4. Can agents hurt each other's feelings? Should they? What about hurting the feelings of humans?

Some of these questions may seem like science fiction – indeed, some of them may *be* science fiction – but again, they can help us to think more deeply about the nature of artificial agents, and about what it means to be human.

The social sciences have accumulated much information about how human societies and human communication work. It is not yet very clear what specific topics may be of greatest relevance to agent research, but I would like to briefly discuss some areas that have attracted my own interest. Before doing so, I would like to say that, despite occasional claims to the contrary, at least in the research areas described below, there is clearly an ongoing accumulation of empirical knowledge and theoretical explanation, much like that found in the physical sciences (though of course with some methodological differences).

Research in the area of ethnomethodology called conversation analysis ("CA" for short, see [30] for many details) has emphasized the key role played in conversation by what they call "recipient design," the phenomenon that speech is carefully crafted to match the needs and capabilities of recipients, taking account of such varied factors as shared background information, shared values, linguistic competence, attention span, shared past history, and much more; the ability to do this is part of what it means to "know" someone, and it is also reciprocal, i.e., you expect speech directed at you to be recipient designed to the extent that you know that the speaker knows you. There are many subtle aspects to the way that this plays out in real time, as we all know from reflection on personal experience with embarrassing situations in which it is especially difficult to know what to say.

Work in CA has also highlighted the incredible accuracy of timing in ordinary conversation (e.g., in turn taking, agreement markers and adjacency pairs [31]), down to a millisecond, which is in sharp contrast to the approximately 500 milliseconds required for conscious intervention. Again, we all know from our own (sometimes unfortunate) experience that talking often goes ahead of thinking.

Requirements engineering, or at least those variants of it that are socially sensitive, could also contribute to the design of effective agent technology,

since a significant feature of the current situation is that we don't really know much about what systems we *should* build. In previously publications I have described requirements engineering as "the reconciliation of technical and social issues" [8, 9], and have also surveyed methods for eliciting requirements [15]. A major lesson here is that the chances of getting a good system are much greater if you have very clear ideas of what it is for and who its users are; agent projects that are technology-driven often neglect this aspect of development. Another lesson is that getting good requirements usually involves careful and often difficult work with real users; the imagination of the designer is almost never an adequate substitute [15].

Distributed cognition is another area with rich implications for socially intelligent agents; here I would particularly mention the work of Edwin Hutchins [18] and Lucy Suchman [35], which carefully describe many important kinds of interaction that are rarely considered for artificial agents, such as mentoring, telling stories, and using plans as resources for coordination rather than action. Moreover, the ways in which (for example) replanning, monitoring other agents, and delegating tasks are handled by humans are often far from what one might expect based on experience with programming.

The area of the sociology of technology and science called **actor-network theory** ("ANT") also has much relevance to agent research. Founded in France by Bruno Latour and Michel Callon [26, 4], it is being applied, criticized, and extended at sites around the world, including UCSD, through work of Leigh Star and Geoffrey Bowker [3], among others. This approach focuses on *networks* of actors, connected by *links*, rather than on autonomous agents; the nodes in these networks include not just human actors, but also non-human actors, which may be artifacts not normally thought of as actors at all, e.g., PCs, programming languages, and transmission media; for this reason, the more neutral word **actant** is often used. Work, consisting of chains of translations along links among actants, must be done to hold the network together, "recruiting" actants to contribute by translating into their languages and values. This is the mechanism by which socio-technical compromise is achieved and through which the work of socio-technical projects gets coordinated [27].

Star and Bowker [3] add to classical ANT an emphasis on the infrastructure required to support networks, on the technical standards and classification systems that allow complex constructions to be accomplished, and on what gets ignored or left out in the stories that are told about projects. To this, my own recent research would add an emphasis on the *values* that are embodied in actants, and that are translated along the links; my claim is that social and ethical concerns are inseparable, and that they are ubiquitous in all socio-technical systems, even though they may be hidden; see the notes for my UCSD class on social and ethical aspects of information technology [12], especially chapters 5 and 6, for more detail on this and related topics. What may be especially relevant is that there are precise sociolinguistic techniques [15] for uncovering these values and the roles that they play.

Narratology is the theory of stories. A surprising amount of precise information is known about the structure of narratives of personal experience, due especially to work of William Labov [21]. In particular, after an optional *orientation section*, a typical narrative consists of a sequence of *narrative clauses* interleaved with *evaluative material*, which places the narrative material in its social context by bringing out shared values that are involved. The **narrative presupposition** says that the order of the narrative clauses reflects the order of the events they report. Much more could be said (e.g., see material in [12]), but this will suffice for present purposes.

Although more cognitive than social, work by Lakoff [22] and others on metaphor has significant implications for socially intelligent agents; this work shows that metaphor, far from being an esoteric literary device, is pervasive in ordinary human interaction. It is no exaggeration to say that metaphors are the most significant tool we have for understanding and communicating in many key areas of life. For example, Johnson [19] shows that metaphors are the foundation for much of our real world moral reasoning, while Lakoff and Nunez [23] show that metaphors are the foundation for much of our real world understanding of mathematics.

A key aspect of using metaphors is to **blend** several metaphors into a single conceptual space, as shown in important work by Gilles Fauconnier and Mark Turner [6, 7, 37]. Blending, also called conceptual integration, is done unconsciously and almost

instantaneously, so that we hardly notice it happening; nevertheless, it is pervasive and important. Without the capability to process complex blends, agents will never achieve a more than superficial understanding of what people really communicate. Computational aspects of blending have been little explored, but see [10] for a formalization that serves as the foundation for an experimental algorithm now being implemented at UCSD.

## 6 Alternatives

The near amicable resolution of the dispute between Shneiderman and Maes suggests that the dichotomy between delegation to agents and direct manipulation was always artificial. First, both of them are metaphors, which can be realized in many different ways, and which typically occur blended with other metaphors, such as information visualization, windows, menus, buttons, and the like. Second, these two are not the only major alternatives; in particular, navigation is an important alternative, which of course can also be blended with other metaphors.

The Meaning and Computation Lab at UCSD has been experimenting with a variant of navigation called **narrative-driven navigation**, in our Kumo system [13]. Kumo can be considered an agent that assists users with proofs, and then generates websites that display those proofs. Kumo can be used either in a stand-alone mode, or in a distributed mode that supports a community of provers cooperating on a project. The second mode is implemented using a specialized protocol that maintains proof consistency over a distributed database [16], with each user delegating to a local copy of Kumo, which knows (subject to the usual internet communication delays) what other users have succeeded or failed to do. In this mode, Kumo can be considered a distributed community of cooperating agents, each under the control of a single user. The output of Kumo is viewed on a web browser, and follows a narrative-driven navigation metaphor, blended with direct manipulation and other metaphors.

The narrative-driven navigation metaphor applies to a user seeking to understand what happened in a given proof attempt (such a user may of course be the same one who produced the proof, and who now wants to know whether it succeeded, and if it failed, wants to know what to try next). Proof websites are structured as narratives, to improve the comprehension of proofs, by utilizing the temporal and evaluative conventions of narrative, including the possibilities of dramatic tension, as advocated by Aristotle [1], who long ago said "drama is conflict."

The agent delegation metaphor applies to the prover, who tells Kumo how to do the proof and how to display it; this may consist of hints and/or detailed commands, and can also rely on defaults. The architecture of Kumo and its interfaces are far from general purpose, and it is arguable whether the use of agent terminology is much help in describing it. See [13, 14] for technical details, and see [10] for the use of algebraic semiotics in the interface design. Generalizing from this experience, we recommend using a carefully chosen blend of interaction metaphors when building interfaces for complex systems.

## 7 Conclusions

As we have explored the debate concerning agents, it has become clearer that we don't know very much about agents, either human or artificial; we have found more questions than answers, and the unanswered questions seem more important than the answers that we currently have. However it does seem safe to predict that agent technology, whether or not it retains that name, will continue to reinvent itself, probably in ways that we cannot imagine at this time. It also seems safe to predict that substantial moral problems will be raised, including some that we cannot now imagine. It is difficult to believe that there will be any technical breakthroughs that permit agents to come anywhere close to being human, but it is safe to bet that we will continue to try, and in the process will learn more about why it is so difficult to implement truly human-like agents, and that many of these difficulties have a social origin. I hope that such experiences may inspire a bit more humility about our technological capabilities, and perhaps even a bit of awe at what it means to be human.

## References


[1] Aristotle. *Poetics.* Dover, 1997. Translation by S.H. Butcher; original from approximately 330 B.C.

[2] Charles Babbage. *The Ninth Bridgewater Treatise: a Fragment.* C. Knight, 1837.

[3] Geoffrey Bowker and Susan Leigh Star. *Sorting Things Out*. MIT, 1999.

[4] Michel Callon. Elements of a sociology of translation. In John Law, editor, *Power, Action and Belief: A New Sociology of Knowledge?*, pages 196–233. Routledge and Kegan Paul, 1986.

[5] Hubert Dreyfus. *What Computers Still Can't Do*. MIT, 1992.

[6] Gilles Fauconnier and Mark Turner. Conceptual integration networks. *Cognitive Science*, 22(2):133–187, 1998.

[7] Gilles Fauconnier and Mark Turner. Principles of conceptual integration. In Jean-Pierre Koenig et al, editor, *Conceptual Structure, Discourse and Language, II*, pages 269–283. CSLI (Cambridge), 1998.

[8] Joseph Goguen. Requirements engineering as the reconciliation of social and technical issues. In Marina Jirotka and Joseph Goguen, editors, *Requirements Engineering: Social and Technical Issues*, pages 165–200. Academic, 1994.

[9] Joseph Goguen. Formality and informality in requirements engineering. In *Proceedings, International Conference on Requirements Engineering*, pages 102–108. IEEE Computer Society, April 1996.

[10] Joseph Goguen. An introduction to algebraic semiotics, with applications to user interface design. In Chrystopher Nehaniv, editor, *Computation for Metaphors, Analogy and Agents*, pages 242–291. Springer, 1999. Lecture Notes in Artificial Intelligence, Volume 1562.

[11] Joseph Goguen. The ethics of databases. In Tara McPherson and Nina Wakeford, editors, *Confronting Convergence*. Erlbaum, to appear.

[12] Joseph Goguen. Social and ethical issues in information technology, Winter 2001. `www.cs.ucsd.edu/users/goguen/courses/175`.

[13] Joseph Goguen and Kai Lin. Web-based multimedia support for distributed cooperative software engineering. In *Proceedings, International Symposium on Multimedia Software Engineering*, pages 25–32. IEEE, 2000. Held in Taipei, Taiwan, December 2000.

[14] Joseph Goguen, Kai Lin, Grigore Roşu, Akira Mori, and Bogdan Warinschi. An overview of the Tatami project. In *Cafe: An Industrial-Strength Algebraic Formal Method*, pages 61–78. Elsevier, 2000.

[15] Joseph Goguen and Charlotte Linde. Techniques for requirements elicitation. In Stephen Fickas and Anthony Finkelstein, editors, *Requirements Engineering '93*, pages 152–164. IEEE, 1993. Reprinted in *Software Requirements Engineering (Second Edition)*, ed. Richard Thayer and Merlin Dorfman, IEEE Computer Society, 1996.

[16] Joseph Goguen and Grigore Roşu. A protocol for distributed cooperative work. In Gheorghe Stefanescu, editor, *Proceedings, FCT'99, Workshop on Distributed Systems (Iaşi, Romania)*, volume 28, pages 1–22. Elsevier, 1999. Electronic Lecture Notes in Theoretical Computer Science.

[17] Martin Heidegger. *The Question Concerning Technology and other Essays*. Harper and Row, 1977. Translations by William Lovitt.

[18] Edwin Hutchins. *Cognition in the Wild*. MIT, 1995.

[19] Mark Johnson. *Moral Imagination: Implications of Cognitive Science for Ethics*. Chicago, 1993.

[20] Immanuel Kant. *Critique of Judgement*. Macmillan, 1974. Translated by J.H. Bernard.

[21] William Labov. The transformation of experience in narrative syntax. In *Language in the Inner City*, pages 354–396. University of Pennsylvania, 1972.

[22] George Lakoff. *Women, Fire and Other Dangerous Things: What categories reveal about the mind*. Chicago, 1987.

[23] George Lakoff and Rafael Núñez. *Where Mathematics Comes from: How the Embodied Mind Brings Mathematics into Being*. Basic Books, 2000.

[24] Dalai Lama. *Ethics for the new millenium*. Riverhead Books, 1999.

[25] Jaron Lanier. Agents of alientation. *Journal of Consciousness Studies*, 2:76–81, 1995.

[26] Bruno Latour. *Science in Action*. Open, 1987.

[27] Bruno Latour. *Aramis, or the Love of Technology*. Harvard, 1996.

[28] Toyoaki Nishida. Workshop on socially intelligent design, 2001. `www.synsophy.go.jp/sid2001`.

[29] Keiji Nishitani. *Religion and Nothingness*. University of California, 1982.

[30] Harvey Sacks. *Lectures on Conversation*. Blackwell, 1992. Edited by Gail Jefferson.

[31] Harvey Sacks, Emanuel Schegloff, and Gail Jefferson. A simplest systematics of the organization of turn-taking in conversation. *Language*, 504:696–735, 1974.

[32] Ben Shneiderman. *Designing the User Interface*. Addison Wesley, 1998. Third edition.

[33] Ben Shneiderman. Humans unite! *Scientific American*, March 1999. Interview by Tim Beardsley.

[34] Ben Shneiderman and Pattie Maes. Direct manipulation vs. agents. *Interactions*, 4(6):43–61, 1997.

[35] Lucy Suchman. *Plans and Situated Actions: The Problem of Human-machine Communication.* Cambridge, 1987.

[36] Kathleen Swigger. Issue focussed on artificial agents. *Computer*, 34(7):39–75, 2001.

[37] Mark Turner. *The Literary Mind.* Oxford, 1997.

[38] Alan Wexelblat and Pattie Maes. Issues for software agent UI. *interactions*, submitted for publication in 1997, to appear.