

Conceptions of the Social that Stand Behind Artificial Intelligence Decision Making

John Monberg

Abstract

AI proponents possessed a seemingly odd predilection to tell stories about times in which no stories are or will be told. Their stories cover a range of time that exceeds that of human experience, beginning with a kind of creation myth about competing songs that are parasitic on the behavior of apes to trajectories of progress in which Man is finally superseded by Machine. AI researchers, funders, and enthusiasts attempt to redefine fundamental social and political concepts of intelligence, meaning, and agency. Their redefinitions emphasize a calculating, controlling, one-dimensional form of rationality, serving to legitimize and extend the power of an already powerful elite. AI theorists ignore the social ground of intelligence, the connection between their computers and the world, and most importantly, the connection between society and their own work. If we accept their claims as true, then their definitions re-order and restructure the social spaces we inhabit.

Introduction

When the 1980s began, computers were not part of the fabric of everyday life for most educated Americans, instead they were understood to be large, expensive mainframe machines requiring specialized facilities and the care of experts. By the end of the decade, personal computers, owned by millions of Americans became a familiar part of the cultural landscape, from Hollywood movies to *New Yorker* cartoons. During this time period artificial intelligence (AI) had matured as an academic discipline. The promises made about the possibility of computer-based intelligence that had been made for decades attracted government funding and media attention, but these promises were unfulfilled as the decade ended. This critical time period offered a chance for reflection about the place of science and technology in the world, in particular a focus on core aspects of intelligence.

To a great extent, the opportunity for reflection about intelligence was lost. This opportunity was foreclosed because the stories that explained and justified the artificial intelligence project were carefully constructed by proponents so that the chaos, uncertainty, and social and environmental complexity built into the deepest core of

AI was left out of their stories. AI proponents possessed a seemingly odd predilection to tell stories about times in which no stories are or will be told. Their stories cover a range of time that exceeds that of human experience, beginning with a kind of creation myth about competing songs that are parasitic on the behavior of apes to trajectories of progress in which Man¹ is finally superseded by Machine (Feigenbaum and McCorduck 1983). Upon careful reading of these stories, a common theme emerges. Through their stories, AI researchers, funders, and enthusiasts attempt to redefine fundamental social and political concepts of intelligence, meaning, and agency. Their redefinitions emphasize a calculating, controlling, one-dimensional form of rationality, serving to legitimize and extend the power of an already powerful elite (Hoffman 1990).

I begin by briefly describing the context in which the AI efforts originated and expanded. The second part of this article explores the social aspects of intelligence and meaning making, aspects which set fundamental limits for any asocial, disembodied AI project. The final section examines the rhetoric of two AI partisans. I critique Marvin Minsky's connectionist form of a Society of the Mind and the Cyc mega-expert system project because they are prominent accounts of the major strands of the AI enterprise.

Ideas arise in a culture and they are shaped by that culture. These ideas in turn, can function to generate political capital, furthering the interests of their proponents. Support can accrue in direct forms, for example, increased levels of funding for specific projects. More importantly, support can be garnered in indirect forms by generating increased legitimacy for a certain type of political order. Ideas expressed as narratives that make sense of, and offer definitions of, the world consequently ought to be considered of central importance. The power of narrative to set the public agenda has been a frequent topic of inquiry in a general political sense (Lasswell 1977; Edelman 1967; Edelman, 1988; Feldman 1989), as well as in a more particular sense for science (Dickson 1988; Nelkin 1987;

Wuthnow 1987; Ezrahi 1990). The formation of a potential common-sense understanding of the world is of prime cultural and political importance because the process of meaning construction is hidden, and people take as “how the world simply is” what may be only in the interest of a narrow elite (Geertz 1983). The creation of persuasive ideologies and systems of meaning grants political power, whether these beliefs spread through the mass media or diffuse through face-to-face interactions. Such power reduces political conflict, encourages the acquiescence of a majority of the population, reduces the space available for critical reflection, and functions as normalizing discourse (Adorno 1990; Hardt 1992; Thompson 1990). In limiting the scope of social imagination, such narratives set the framework for all decisions made about the funding levels, goals, priorities, and expectations for AI technologies.

The narratives surrounding AI are important because the computer is such a powerful metaphor in our society. Computers are a defining technology as we think about human capabilities, agency, and our place in the world. When rights and responsibilities are framed in terms of the computer, these conceptions have direct political repercussions. The popular literature burgeons with examples like the account in *Scientific American* that begins by stating bluntly, “The brain is a remarkable computer” (Hinton, 1992). We are redefined as information processors in a world that is held to be an environment of information to be processed. “Thus, human beings and computers are two members of a larger class defined as information processors, a class that includes many other information-processing systems – economic, political, planetary – and, in its generality, a class that threatens to embrace the universe” (McCorduck, 1988, p.74). This threatening embrace may turn out to be not merely metaphorical when information systems mediate global decision making in fields as consequential as military force projection, flows of financial investments, and environmental monitoring and modeling.

The AI literature continues the long tradition of epistemological certainty and self-righteousness exemplified by Descartes, Hume, Bertrand Russell and the logical positivists. A pointed aggressiveness appears time and time again in the rhetoric of AI practitioners. All previous modes of knowledge that cannot be readily assimilable to AI forms are no longer valid.

They simply are no longer worth knowing. If the position of the most vigorous AI proponents is taken seriously as a model for human agency, some fear we may fall into a rationalized, closed system in which Weber’s iron cage of bureaucracy reaches full fruition and from which there might be no escape;

The increase of computer use in society and in all scientific disciplines could lead to an unforeseen consequence: the impossibility of thinking outside the dominant paradigm. The paradigm of computer culture would become part of the culture, if not all. A troublesome techno-culture of calculus, where policy has no meaning anymore since it is supported by so-called clarified criteria; where alternatives are also ranked by supposedly less enigmatic and erratic procedures; because computing has become ‘laws of thought.’ (Berleur 1990, p. 415)

The AI community constitutes one branch of a broader worldview. This worldview understands technology as a new type of cultural system that restructures the entire social world as an object of social control. This worldview has in turn provoked a rich tradition of social analytical critique. In the perspective of these critics, technology, either inherently or as a tool of elite control, generates domination in the social and natural worlds (Ellul 1964; Merchant 1980; Habermas 1987). Analysts have explored the alienating and repressive role of technology in the workplace (Braverman 1974; Weizenbaum 1976; Noble 1984; Feenberg 1991). A growing literature examines the potential or actual uses of information technology in particular, to effect a more stringent degree of control in the workplace (Clement 1988 1990; Roszak 1986).

Sophisticated, capital-intensive technologies are not developed in a social vacuum, but are developed to meet the needs and further the goals of the groups that fund them. Support for the AI community has come primarily from military, and to a lesser extent, corporate sources. Justification for this largely public funding has been framed in terms of military force projection and multiplication, and corporate productivity and competitiveness (especially after the establishment of Japan’s much-ballyhooed Fifth Generation Project).

The Defense Department’s Advanced Research Project Agency (DARPA) has been a

prime supporter of AI projects, establishing the Strategic Computing Initiative in the pursuit of voice recognition, machine vision, and battle management for the Strategic Defense Initiative. The hundreds of millions of dollars channeled by this organization has been integral to the establishment of every major AI research community; those at the RAND Corporation, MIT's Lincoln Laboratories, Carnegie-Mellon University, the Stanford Research Institute, and the consulting group Bolt, Beranek and Newman (Johnson, 1986, p. 129). Defense dollars supported the work of virtually every light in the AI pantheon: John von Neumann, Herbert Simon, John McCarthy, Alan Turing, Allen Newell, and Marvin Minsky (Minsky, 1985, pp. 323-324). In short, AI is a product of military funding. It is then not surprising that so much of the work done in AI assumes a mechanistic universe, an overly-narrow rationality governed by formalizable and programmable rules, a sense of objective knowledge that proceeds with a neutral, universal logic uncontaminated by social and political "impurities," and an emphasis on refinement of technique and information technology as an instrument of administration in pursuit of more precise control over the natural and the social world. AI researchers manifest a common blind spot with regard to their own work: they see themselves in a quest for "disinterested," "universal," and "value-free" knowledge which supports an endeavor which is nothing if not supremely interested, value-laden, and politically potent.

Social Ground of Intelligence

Having begun with rapid progress in the mid-1950s—an early example was Newell and Simon's General Problem Solver—AI practitioners made bold predictions that the possibility of understanding the universal logic of intelligence would soon be within reach. Almost four decades later, the AI project has made little progress toward reaching its ultimate objective. This lack of progress has not been due to a lack of funding, or to accidental circumstances. The AI project has failed to progress as expected because, as it has been carried out up to this point, the AI project has assumed an impoverished model of intelligence, a model subject to strict and inherent limitations.

With close ties to psychology and analytic philosophy, the AI project assumed that intelligence is located within independent, atomistic individuals; that humans are Cartesian knowers

in fundamental respects. The aspects of intelligence stemming from the complex interactions of embodied, social, experiential, and cultural learners and doers have been virtually ignored. Social cognition is "a domain about which cognitive science and the attendant philosophical literature have had virtually nothing to say" (Jackendoff, 1991, p. 420). This impoverished conception of self has come under increasing attack from a broad range of phenomenologists, hermeneuticists, feminists, pragmatists, and other intellectual camps. Common to these groups is the belief that the self is not an isolated being and can only be understood as an actor in a social context. This richer conception of self and of intelligence has been taken up recently in a variety of ways, in disciplines including psychology (Hermans, 1992) and political theory (Dallmyr, 1984), to offer only a few of the many possible examples.

The interactive, social understanding of self derives in part from the work of George Herbert Mead. The Meadian concept of mind requires the ability to take the point of view of another, requiring from the outset an understanding of the social dimensions of self as "selves can only exist in definite relations to other selves" (Mead, 1963, p. 46). Even the possibility of becoming a social self requires interaction with another social self. As the self only comes into existence as a social being, the interactive aspects of self are central to an analysis of intelligence. As Marcelo Dascal noted, "It is not by digging deeper into the individual's head that one discovers the relevant parameters of his mental life. For these parameters are social, not individual, public, not private, context-relative, not universal" (1989, p. 40). It is only through an analysis of the social, that mystifying claims can be avoided. Dascal continues, "It is only by reference to such a context that these allegedly 'mental' phenomena can be understood and accounted for in a non-mysterious way" (1989, p. 42).

That intelligence only emerges in a social setting holds for computers as well as humans. Without the experience of social life, computers cannot be understood as intelligent creations:

Computers are not part of the social process; they are not personalities for whom a life process is a unity of biological, psychological and social processes. In order to understand meanings or meta-meanings in the context of communication between

human beings, and in order to form relevant social values, the computer must have lived a practical life which is changing the world in sensory, concrete terms. (Fuchs-Kittowski, 1990, p. 465)

AI researchers have failed to pursue the creation of “social” computers. The programs created attempt to distance the AI system not only from the social world, but also from any connection to the world outside the system. AI theories create representations of the world, not connections to or interactions with the world. Jerry Fodor terms this, “methodological solipsism”—“the machine lives in an entirely notational world; all its beliefs are false” (1981, p. 315).

What should count as an example of intelligence? It is clear that a rote enactment of preset rules, regardless of circumstances, does not qualify as an exhibition of intelligent behavior. We do not grant the microwave oven a robust sense of intelligence. The minimal requirement for intelligence is sensitivity to the surroundings of a creature. Thus a computer must not merely respond inflexibly to a stimulus in an environment, but must possess the ability to respond appropriately to a variety of possible situations. In a rule-driven system, rules must be carefully optimized to respond to a specific kind of environment. The system requires different sets of rules to handle different types of circumstances. If the system is completely rule-driven, then a high-level set of rules must determine which set of lower-level rules to execute. But in order to guide this higher-level set of rules in a flexible manner, there must be a still higher level of rules. This never-ending hierarchy falls victim to the Wittgensteinian regress. Implicit in every case of rule following is a *ceteris paribus* condition regarding the application of the rule that cannot be understood within the terms of the rule specified (Dreyfus, 1979, pp. 56-57):

In social life, rules and language games are always embedded in practice, and this practice bridges the gap between rules and their application. Through interaction, participants in a conversation can negotiate an understanding without being trapped within an infinite loop. Wittgenstein’s arguments against the possibility of a private language also apply to any AI approach which functions as a closed system. Because the system is closed there exists no possibility of a check, the outside world has no purchase on

the interpretation of the system. Truth verification for such a system would be the equivalent of reading copy after copy of the same newspaper to verify facts read in the first edition. This limitation sets a boundary condition for a closed AI system, it is necessary to take into account the context as a factor which is capable of changing completely the initial semantic interpretation.

Given the paradoxes of rule following, how is it that human beings can be considered intelligent? We are not trapped within the constraints of a formal system. We interact with others in settings that are open in important respects, creating gaps between our beliefs and our experience of the world. The everyday world provides the backdrop that an analysis of formal rules can never provide, making it possible to act in a contingent world, rule-governed creatures without a theory of action, without always already understanding what all possible rules are. We are able to accomplish contextual definitions of the circumstances we find ourselves in using culturally available conventions.

people have told each other stories and listened to stories in all cultures at all times. In doing so, people arrive at an understanding and ordering of the world and the self. (Hermans, 1992, p. 23)

AI adherents continue to tell stories about the world, stories which aim towards an understanding and an ordering of the world. They tell stories about the possibility of a rule-based, acontextual intelligence, stories about the overcoming of stories. Two of the most important stories are told by Marvin Minsky in his *The Society of Mind* and by supporters of the Cyc mega-expert system.

What is mind? Minsky’s model of mind is the corporate bureaucracy, a metaphor that Minsky returns to time and time again in his book. Wired into the brain is a tiny, finely-organized, complex corporation. For Minsky the homology is almost perfect – the mind becomes a corporation *par excellence*. The title of his book is misleading, for he focuses on activity internal to a disembodied mind, with little connection to the outside world, and refers to the social world only in passing. In Minsky’s mind, a hierarchical structure of “subordinate” agents pass information up and down a management chain directed by “boss” agents at the top of the

pyramid. The discussions of cross linkages of agents at the same level are few.

Minsky is obsessed with control. He envisions a most intricately adjusted system of rewards and penalties that has evolved to ensure that every subordinate part functions according to plan, or will be made to do so in short order. Although the snippets taken from St. John, Shakespeare, and Simone DeBeauvoir that are sprinkled throughout Minsky's writing might indicate a humanist sensibility, Minsky believes every facet of our cultural heritage functions as a control system.

In Minsky's Mind, language functions as a system of control, "If we're to understand how language works, we must discard the usual view that words denote, or represent, or designate; instead, their function is control: each word makes various agents aware of what various other agents do" (p. 196). As do emotions, "Our earliest emotions are built-in processes in which inborn proto-specialists control what happens in our brains" (p. 172). And so for social institutions in general, "All human organizations evolve institutions of law, religion and philosophy, and these institutions both adopt specific answers to circular questions and establish authority-schemes to indoctrinate people with those beliefs" (p. 49).²

Minsky finds the need for control outside the mind as well, "Those lower-level agents need to be controlled. It's much the same in human affairs. When any enterprise becomes too complex and large for one person to do, we construct organizations in which certain agents are concerned not with the final result, but only with what some other agents do" (p. 34). He appreciates servants who possess the least voice, "No supervisor can know everything that all its agents do. . . The best subordinates are those that work most quietly" (p. 60). This emphasis on control is evident in the structure of his book as well. Minsky has stated that the interconnections between his essays are so varied and complex that a standard book format would be ineffective and that the format of his book itself had to be modified. However these varied connections do not become apparent in his book; on the contrary his format does not encourage flexibility. His book consists of a series of essays carefully arranged and finely categorized in a linear hierarchy ordered from 1.1 to 31.8:

Despite all of his emphasis on control, Minsky neglects to discuss the role of power in his systems. Conflict is interpreted only in the context of miscommunication or lack of information. What limits are to be placed on a control system? What application of control is allowable? What type of power even is to be preferred? How can elements of the system be made accountable to other elements within the system or to larger elements outside of the system? Does it make sense to speak of any range of freedom or autonomy to be accorded to the agents (whether of high or low level) within Minsky's system?

A discussion of power is not all that is missing from Minsky's Mind. Intentionality, understanding, purpose, autonomy, feelings, aspirations are all reduced to a one-dimensional focus on order and control. In his Mind, rules are not only regulative, but also constitutive of the "experience" of the system. What would it be like to be such a system? What would such a system be willing to die for? Or, for that matter, what would such a system have to live for? Minsky's Mind is not a von Neumann machine, a kind of computer architecture in which one central processor executes a rule-driven algorithm. Instead sets of connections exist on many processors that execute instructions concurrently. These connections are not made explicitly by a system programmer, and indeed may not even be understandable to an observer outside of the machine. Connections evolve through trial-and-error processes that reconfigure the connections in ways aimed at minimizing the discrepancy between data input into the machine and the desired output.³

Connectionist machines function differently than explicitly rule-guided machines do, but they remain bound by similar restrictions with regard to the implementation of intelligence. The agents that Minsky envisions in his system cannot have intelligence, or each agent would be a *homunculus*, and Minsky would be assuming at a lower level just what he is trying to enact at a higher level. Minsky is clear on this point, "Each mental agent by itself can only do some simple things that needs no mind or thought at all" (p. 19).

If these agents do not possess even a rudimentary intelligence, his analogy between his Society of Mind and the society of humans breaks down. Subject to the constraints of their

culture, individuals possess a sense of agency in that they can shift their attention between a larger social whole and the parts they play, they can negotiate an understanding of a situation, and they can give accounts of their actions. Minsky's agents are necessarily devoid of agency. They have no sense of their pasts, and no way of taking into account the new or unexpected. Minsky is also clear about this, "Those tiny mental agents simply cannot know enough to be able to negotiate with one another or to find effective ways to adjust to each other's interference" (p. 33). The lack of even a limited possibility of autonomy on the part of these sub-systems in turn sets a limit on the potential the system has for intelligence. The results generated from the connectionist computer systems that have been implemented so far have remained scant. Systems are able to model only "toy problems" (Papert, 1990, p. 13). Even given the great strides in processor speed, the number of interconnections possible, and memory size which have been made in the past decade, the state of the art has not advanced beyond what was possible when computer systems were much less powerful. The possibility of scaling-up these systems to tackle more realistic problems does not seem amenable, even given what will surely be very large advances in computer hardware.

Get A Bigger Hammer Approach: The Mega-Expert System

The Cyc project to create the largest expert system ever constructed began in 1984 at the Microelectronics and Computer Technology Corporation, a research consortium in Austin, Texas. The consortium exists on the support of Apple, Digital Equipment Corporation, Eastman Kodak Corporation, NCR Corporation and other large computer manufacturers and users. The goal of this project is unlike that of most traditional expert system approaches which focus on gathering a great deal of specialized information about a narrowly focused technical area – the most efficient way to deploy a telephone switching network, for example.

Instead, the Cyc system is an attempt to collect and code common-sense reasoning. A person cannot walk through a wall. Water falls downhill. All animals live, die and stay dead. The goal of the knowledge base is the support of 100 million of these common-sense assertions, creating a system that is 10,000 times as large as an average expert system (Harrar, 1990, p. F7).

The rhetoric surrounding the development of expert systems, even the systems of such vast complexity as Cyc, has been more restrained. In this context one does not find the grandiose claims of a universal model of intelligence like those of Simon, Minsky, et al. The claims made on the behalf of these systems are straightforward: they are designed to effect a transfer of the control of knowledge from workers to management. Workers, it is hoped, may be made more cheap, reliable, and productive. Joseph Scullion, director of strategic planning at NCR, explained that, "Just being able to capture common-sense intelligence in a work-station means that whatever application you can run can be more complex. And lightly trained people can be made more productive" (Harrar, 1990, p. F7). Even given the vast quantity of data stuffed into this system, the constraints suffered by any formal rule-based system hold. It is difficult to describe all the relevant attributes of a given context, if it is not known in advance what the criteria for relevance is, or how the criteria for relevance may change over time. Even if the context can be defined appropriately, how will the machine determine when the context changes and what the optimal procedure to follow is? Marcelo Dascal asks, "A system cannot always use the same script or schema. If it is to not behave stupidly, it must be able to shift from one schema to another when required. But how is a system to know when this is required?" (1989, p. 46).

We arrive again at the Wittgensteinian infinite regress; and in the everyday world in which we inhabit, it is no trivial matter to come across exceptions to rules, exceptions to the exceptions of rules, unforeseen circumstances, dashed expectations, new appreciations of old situations, *ad infinitum*. Bruno Latour's (1992) description of doors and other mundane objects provides a rich illustration of the often-unnoticed complexities in our lives that a rigid, formal system would adapt to with great difficulty.

The Moral of the Story

The AI project should be understood as a typical extension of the long Western quest for a kind of universalistic epistemological certainty. AI theorists ignore the social ground of intelligence, the connection between their computers and the world, and most importantly, the connection between society and their own work. This purposeful ignorance allows the AI community to discredit every other form of knowledge, which is replaced by a technocratic, controlling, bureaucratic understanding of the world. If we

accept their claims as true, then their definitions re-order and restructure the social spaces we inhabit. Our world becomes a bit colder and grayer, and our understanding of our own place in the world becomes constricted and diminished.

Dr. John Monberg is an assistant professor in the Department of Communication Studies at the University of Kansas, Lawrence.

References

- Adorno, T. (1990). Culture industry reconsidered. In J. C. Alexander & S. Seidman (Eds.), *Culture and society: Contemporary debates* (pp. 275-282). New York: Cambridge University Press.
- Berleur, J. (1990). Recent technical developments: Attitudes and paradigms. In J. Berleur, A. Clement, R. Sizer & D. Whithouse (Eds.), *The information society: Evolving landscapes*. North York, Ontario: Springer-Verlag.
- Braverman, H. (1974). *Labor and monopoly capitalism*. New York: Monthly Review.
- Clement, A. (1990). Office automation and technical control of information workers. In J. Berleur, A. Clement, R. Sizer & D. Whithouse (Eds.), *The Information society: Evolving landscapes*. North York, Ontario: Springer-Verlag.
- Dallmyr, F. (1984). *Polis and praxis*. Cambridge: MIT Press.
- Dascal, M. (1989). Artificial intelligence and philosophy: The knowledge of representation. *Systems Research*, 6(1), 39-52.
- Dickson, D. (1988). *The new politics of science*. Chicago: University of Chicago Press.
- Dreyfus, H. (1979). *What computers can't do: The limits of artificial intelligence*. New York: Harper.
- Edelman, M. (1967). *The symbolic uses of politics*. Urbana: University of Illinois Press.
- Edelman, M. (1988). *Constructing the political spectacle*. Chicago: University of Chicago Press.
- Ellul, J. (1964). *The technological society*. trans. J. Wilkinson. New York: Vintage.
- Ezrahi, Y. (1990). *The descent of Icarus: Science and the transformation of contemporary democracy*. Cambridge: Harvard University Press.
- Feenberg, A. (1991). *Critical theory of technology*. New York: Oxford Press.
- Feigenbaum, E., & McCorduck, P. (1983). *The fifth generation*. Reading, Mass.: Addison-Wesley.
- Feldman, M. (1989). *Order without design: Information production and policy making*. Stanford: Stanford University Press.
- Fodor, J. (1981). Methodological solipsism. In J. Haugeland (Ed.), *Mind Design*. (pp. 67-94). Montgomery, Vermont: Bradford.
- Fuchs-Kittowski, K. (1990). Information and human mind. In J. Berleur, A. Clement, R. Sizer & D. Whithouse (Eds.), *The Information society: Evolving landscapes*. North York, Ontario: Springer-Verlag.
- Geertz, C. (1983). *Local knowledge: Further essays in interpretive anthropology*. New York: Basic Books.
- Habermas, J. (1987). *The theory of communicative action: Volume two, lifeworld and system: a critique of functionalist reason*. Boston: Beacon Press.
- Hardt, H. (1992). *Critical communication studies communication, history and theory in America*. London: Routledge.
- Harrar, G. (1990). The software with good sense. *The New York Times*. April 1, F7.
- Hermans, H. (1992). The dialogical self: Beyond individualism and rationalism. *American Psychologist*, 47, 23-33.
- Hinton, G. (1992). How neural networks learn from experience. *Scientific American*, 267 (September), 144-151.

- Hoffman, R. (1990). John McCarthy: Approaches to artificial intelligence. *IEEE Expert*, 5(3), 87-89.
- Jackendoff, R. (1991). The problem of reality. *Nous.*, 25, 411-433.
- Johnson, G. (1986). *Machinery of the mind*. Redmond, Washington: Tempus.
- Lasswell, H. (1977). *On political sociology*. Chicago: University of Chicago Press.
- Latour, B. (1992). Where are the missing masses? The sociology of a few mundane artifacts? In W. Bijker & J. Law (Eds.), *Shaping technology/building society* (pp. 225-258). Cambridge: MIT Press.
- MacIntyre, A. (1984). *After virtue*. Notre Dame: University of Notre Dame Press.
- McCorduck, P. (1988). Artificial intelligence: An apercu. In S. R. Graubard (Ed.), *The artificial intelligence debate* (pp. 65-83). Cambridge: MIT Press.
- Mead, G. H. (1963). *Mind, self, and society*. Chicago: University of Chicago Press.
- Merchant, C. (1980). *The death of nature*. San Francisco: Harper and Row.
- Minsky, M. (1985). *The society of mind*. New York: Touchstone.
- Nelkin, D. (1987). *Selling science*. San Francisco: W. H. Freeman and Company.
- Noble, D. (1984). *Forces of production*. New York: Oxford University Press.
- Oakeshott, M. (1991). *Rationalism in politics and other essays*. Indianapolis: Liberty Press.
- Papert, S. (1988). One ai or many? In S. R. Graubard (Ed.) *The artificial intelligence debate* (pp. 1-14). Cambridge: MIT Press.
- Roszak, T. (1986). *The cult of information*. New York: Pantheon.
- Thompson, E. P. (1966). *The making of the English working class*. New York: Vintage.
- Thompson, J. (1990). *Ideology and modern culture*. Stanford: Stanford University Press.
- Weizenbaum, J. (1976). *Computer power and human reason*. San Francisco: W. H. Freeman.
- Wuthnow, R. (1987). *Meaning and moral order: Explorations in cultural analysis*. Berkeley: University of California Press.

Notes

- ¹ The use of the gender-loaded term Man reflects the gender imbalance of the AI research community. More importantly for this essay, it expresses a universalistic, asocial, disembodied model of rationality.
- ² Minsky's emphasis on control is even more pronounced when contrasted with the role tradition plays in lived experience found in the work of authors as varied as Geertz, 1983; Thompson, 1966; Oakeshott, 1991 and MacIntyre, 1984.
- ³ For explanations of the various connectionist approaches see Cowan and Sharp, 1990; Hinton, 1992; and Papert, 1990.

