Artificial Decision-Making and Artificial Ethics: A Management Concern

Omar E. M. Khalil

ABSTRACT. Expert systems are knowledge-based information systems which are expected to have human attributes in order to replicate human capacity in ethical decision making. An expert system functions by virtue of its information, its inferential rules, and its decision criteria, each of which may be problematic. This paper addresses three basic reasons for ethical concern when using the currently available expert systems in a decisions-making capacity. These reasons are (1) expert systems' lack of human intelligence, (2) expert systems' lack of emotions and values, and (3) expert systems' possible incorporation of intentional or accidental bias. For these reasons artificial ethics seems to be science fiction. Consequently, expert systems should be used only in an advising capacity and managers should not absolve themselves from legal and ethical responsibility when using expert systems in decision making.

Introduction

Artificial Intelligence (AI) is the study of ways of using computer based systems to perform tasks or to solve problems normally performed by humans. The development of such systems, which mimic human intelligence, has been a major interest of the AI researchers. AI research encompasses a number of research subfields such as expert systems, robotics, natural languages, and simulation of human sensory capabilities.

Expert systems (ES), as a subset of AI projects, includes both computer technology and knowledge that attempts to achieve expert-level results in problem solving and decision making (Mykytyn *et al.*, 1990). The general recognition of AI as a practical tool has led to the rapid deployment of expert systems by corporate America (Ansari and Modarress, 1990). These systems are either operations or management information systems depending on whether they are being used to give expert advice to control operational processes or to help managerial end users make decisions (O'Brien, 1990, p. 365).

The technological excitement of expert systems, however, must be balanced against the possible social, political, and other organizational implications of these systems. Since ethics deal with conduct (Ferrell and Fraedrich, 1991, p. 35) and corporate decision making is inherently based on the conduct of the decision maker(s), it can be assumed that ethical or unethical behavior is part of the corporate environment. Furthermore, since information systems (IS) is a major part of the corporate decision making process, the deployment of AI and expert systems has expanded the scope of the ethical concerns of information systems.

Research on ethics and information systems in general has addressed a number of issues. Among them are: the social implications of ethics and information systems (i.e., Paradice and Dejoie, 1991; Waldrop, 1987; LaChart, 1986; Mason, 1986; McCorduck, 1979; Boden, 1977; Weizenbaum, 1976), the ethical issues that pertain to identifiable groups of individuals (i.e., Shim and Taylor, 1988; Parker, 1979); and the effects of technology on ethical decision making (i.e., Taylor and Davis, 1989; Bommer *et al.*, 1987; Jastrow, 1987; Sheridan, 1987).

On the other hand, the research that has dealt directly with the ethical issues concerning AI may fall into two broad categories: legal and philosophical (Dejoie *et al.*, 1991, p. 226). Legal issues include the classical ideas of liability and ownership (i.e., Mason, 1986). Philosophical issues refer to more

Omar E. M. Khalil is a visiting Lecturer of information systems at the University of Massachussets at Dartmouth. He has published in the International Journal of Man-Machine Studies *and in several proceedings.*

Journal of Business Ethics **12**: 313–321, 1993.

thought-provoking ideas: the degree of responsibility given to a machine for its role in the decision making process; the evolution of an AI system with rights; the question of experimental AI research, and the question of conscious AI entities (i.e., LaChart, 1986; McCorduck, 1979; Boden, 1977; Weizenbaum, 1976).

While the relevant prior research has addressed a wide range of information systems and AI-related ethical issues, the ethical implications of the "artificial decision making" has received little attention. If it is possible for expert systems to replace human decision makers, the question becomes whether they should be used in such a capacity. Simon (1990) further explains this issue:

> Going beyond artificial intelligence, I think we have "artificial ethics" all over the place today. Every time a computer makes a decision . . . it is implementing a set of goals or values [and] you won't be happy with a program unless you are satisfied with the balance of values that is implemented by the program (p. 666).

The issue is not whether expert systems should be used but in what capacity are they to be used. This paper addresses three of the expert systems' limitations that could raise ethical questions and, consequently, management concerns when using these systems in a decision-making capacity in organizations.

Ethical decision making: background

Ethics is generally viewed as "inquiry into the nature and grounds of morality where the term morality is taken to mean moral judgments, standards and rules of conduct" (Taylor, 1975, p. 1). An ethical issue is a problem, situation, or opportunity that requires the decision maker to choose among several actions that must be evaluated as right or wrong, ethical or unethical (Ferrell and Raedrich, 1991, p. 35). Additionally, business ethics comprises those moral principles and standards that guide behavior in dealing with ethical issues in business organizations.

By their nature, ethical problems involve conflicts of values. Most of the decisions in business organizations involve some degree of ethical judgment and, therefore, may not be taken solely on the basis of arithmetic (Cadbury, 1987). Decision makers should think about the multitudinous frames of reference and the many consequences of decisions before making them.

A variety of models have emerged which outline the major steps or functions involved in the decision making process. Most of the models, however, include the following steps in one form or another: (1) setting objectives; (2) searching for alternatives; (3) evaluating alternatives; (4) choosing an alternative; (5) implementing the decision; and (6) controlling the results. Each of these steps involves the gathering and processing of information within a value construct and the cognitive limitations of the decision maker (Bommer *et al.*, 1987).

At each step in the decision-making process, as Bommer *et al.* (1987) explain, the manager acquires and processes a myriad of information. Some of the information is problem specific, whereas other information relates to environmental factors such as social, governmental/legal, work, professional, and personal. This information ranges from hard data – such as laws and stated corporate policies – to soft data – such as an individual's self concept – with a range of information in between these states. The manager must then synthesize and analyze this information to determine a rational decision to the problem situation.

Ethical issues may arise at any step of the decision making process. In setting objectives, for instance, it is necessary to consider ethical concerns relating to the choice of pursuing various directions. In comparing various alternatives, ethical considerations often arise as part of the valuation process. In the implementation step, the potential consequences to resources – human and physical – which will be affected by the decision, must be considered from the ethical perspective (Bommer *et al.*, 1987). It is also suggested that human information processing – how people gather and use information – is related to the decision-making process (Fleming, 1985). Each decision maker brings a particular information processing approach, and thus decision style, to bear on a particular problem.

Making ethical decisions may be a relatively easy task when all the relevant facts are clear and the choices are black and white. It is a different story, however, when the situation is ambiguous, information is incomplete, and points of view are in conflict. In such situations ethical decisions depend on both the decision-making process itself and on the experience, intelligence, and integrity of the decision maker (Andrews, 1989).

Expert systems, as knowledge-based information systems, attempt to achieve expert-level results in problem solving. They are expected to have human attributes in order to replicate human capacity in ethical or moral decision making. The question, however, is how many human attributes are incorporated in the currently available expert systems?

Expert systems and human attributes

Drawing on the idea of the "moral judge" or the "ideal observer," the attributes of a moral decision maker include: (1) knowledge of all relevant facts; (2) in-biasedness; (3) freedom from disturbing passion; and (4) the ability to vividly imagine the feelings and circumstances of the parties involved (Firth, 1952).

The attributes of the moral judge can be viewed as the conditions under which a valid moral judgment might be made by a decision maker. Considering the first three attributes of the moral judge, a case might be made that an unemotional expert system could be considered a better decision maker than a human expert. Such a system might be able to store and retrieve more factual data and not be disturbed by violent passions and interests (LaChat, 1986). The last attribute, however, is problematic since emotions would be involved.

While LaChat (1986) questions expert systems' emotions and reasoning capability in having the capacity for making ethical decisions, it can be argued also that expert systems' intelligence and in-biasedness are problematic when making such decisions. These issues of expert systems' intelligence, emotions and values, and in-biasedness are discussed in the following pages.

Intelligence

The common practice of building expert systems is to have the knowledge engineer, together with the human expert, solicit and write the rules that represent the expert's thought process. Human experts, however, don't think in terms of fixed values, but they think in terms of real world values (Cox and Goetz, 1991).

Experts are seldom able to retrace that analytic steps followed to make a particular decision. They may be able to highlight important factors that went into a decision but usually cannot describe the whole process. In the process of becoming expert, individuals distill many observations into intuition and deep understanding but cannot articulate them (Dreyfus and Dreyfus, 1986, pp. 193—201). That fact makes it hard to mine the expert's knowledge (Leonard-Barton and Sviokla, 1988).

Walter (1988) further explains the difficulty of mining the expert's brain: ". . . it will be impossible to diagnose a not-observable, not-controllable, real-time 'computer' whose detailed blueprint is not known, and whose components have complex dynamics not fully known, and are packed tighter than will allow a probe between them . . ." LaChat (1986) adds that mapping the thought patterns of a human expert would have to include memories and experiences as well as hopes, aspirations, and goals. Such a map, which would have to include the three temporal dimensions of human consciousness – past, present, and future – is too complex to be programmed.

Weizenbaum (1976, p. 223) also argues that the unconscious aspect of the human mind cannot be explained by the information-processing primitives. It is wrong to assert that any scientific account of the whole man is possible. The man and machine are separated by qualities – courage, trust, risk, endurance, and tenacity – that are applicable to one and not the other. The problem in such cases is that some designers may build into their programs routines that give the impression of qualities such as those mentioned (Newman, 1988).

More often than not, all the knowledge used by a human decision maker cannot be captured in the development of knowledge-based expert systems. Consequently, most expert systems are built for narrowly specified domains of knowledge (Liebowitz, 1987). They are often designed to help with only one component of a task, and the human user is left to interpret the output and buffer it from problems beyond its range (Sheil, 1987).

Thus, there are programs, for example, that are expert on infectious diseases but know nothing of general medicine. So, an expert system designed to diagnose heart disease could possibly make intelligent-sounding but completely misguided recommendations for a patient with a broken leg. The danger is that users will mistake the intelligent tone for real competence and act on the machine's advice (Sheil, 1987).

There is a distinction between the specialized expert system's intelligence and a human expert's intelligence. Although it is obvious that machines can perform some activities at a higher level than persons can, these tasks remain, by and large, highly specialized and therefore remote from the capacity of human intelligence for multipurpose activities (LaChat, 1986). In other words, a personal intelligence must have a personality that is formed through time and through lived experience and must certainly include the emotional.

Consequently, if an overall level of cognitive ability comparable to the fall range of human intelligence is expected from Artificial Intelligence projects, then expert systems are not considered to be intelligent. Whether the ongoing Artificial Intelligence research program could possibly solve the problem of conscious intelligence and of producing systems that think is at best uncertain (i.e., Rosenfield, 1991; Churchland and Churchland, 1990; Searle, 1990; Lloyd, 1985); and this possibility seems to be contingent on both the availability of valid cognitive models of human intelligence and the successful implementation and manipulation of these models in computers.

The inadequacy of "human" intelligence in expert systems limits their use in a decision-making capacity. This limitation is attributed to the following factors: (1) in decision-making situations, sometimes pre-defined objective criteria can be applied to a task, but often human values must be applied and flexibility in decision-making is a necessity (Weizenbaum, 1976); (2) the ability to have deep-reasoning systems where expert systems are constructed for general functional areas like classification or diagnosis is essential, but most expert systems today do not exhibit this quality; and (3) expert systems cannot learn from their mistakes (i.e., if the user gives the expert system the same input as when a mistake in advice was made, the expert system will make the same error.) Work is being done in learning and common sense reasoning to overcome this problem, but a great deal of research is needed to solve this limitation (Liebowitz, 1987).

Given the inadequacy of expert systems' intelligence, expert systems become useful only when the decision problem is well defined (Newman, 1988; Barnett, 1982). In addition, decision making situations are not homogeneous. Therefore, the decision maker should be able to move from the general to the specific; from the ill-defined to the highly structured problems or from high-level to low-level forms of analysis. The human decision maker has an acquaintance with the real world that cannot be taken for granted in a computer program (Blois, 1980). In such heterogeneous decision situations, specialized expert systems tend to find their natural use as assistants to humam experts.

Emotions and values

Many look upon AI research with fear that it might result in the creation of Frankenstein's monster (Lloyd, 1985). Perhaps that fear was the driving force for Asimov's (1964) imagination of what he calls the "rational robot." The artificial brain of such a rational robot would be constructed with three safeguard rules:

(1) A robot may not injure a human being or, through inaction, allow a human being to come to harm.
(2) A robot must obey the orders given it by human beings except where such orders would conflict with the first rule.
(3) A robot must protect its own existence as long as such protection does not conflict with the first or second rule.

Asimov's rational robot – like Firth's (1952) moral judge – is an imaginative AI entity with built-in artificial ethics. While Asimov's rules can be viewed as standards or guidelines for building intelligent machines in general, the first law is crucial in building expert systems in particular. That is, the outcomes of expert systems, whether they are advices or decisions, should not wrongly harm a human being.

The concern about the use of expert systems, however, is due to the fact that these systems are not equipped with any built-in ethical system analogous to Asimov's three rules of robotics (Waldrop, 1987; p. 242). Also, a computer-based system only does what it is programmed to do, and it can never become moral or ethical without the mysterious property of free will (LaChat, 1986). In addition, an expert system which lacks the capacity to value and to make moral choices cannot be considered a moral judge, no matter how intelligent it is.

The goal of AI researchers in general, and expert systems in particular, has been construed as the reduction of the real world's complexity into simple calculations that can be easily implemented. Reductionism, however, can never construct more than a shallow reflection as well as the tendency to devaluate those parts that cannot be quantified (Newman, 1988; Dreyfus and Dreyfus, 1986, pp. ix–xv). The problem will be even worse, Gould (1985) argues, when taking imprecise language that expresses uncertain and changing knowledge, and then to harden it into mechanical devices. This process is characterized as making a move from the human world, where thinking is possible, to the mechanical world, where thinking cannot exist. Consequently, the question of whether all thinking can be formalized in some sort of rule structure becomes a crucial one for the development of expert systems (Churchland and Churchland, 1990; Negroponte, 1990; Searle, 1990; Waldrop, 1988).

Consider, for instance, that a knowledge engineer attempts to develop a program to capture the medical diagnostic ability of a certain physician. The knowledge engineer tries to break the physician's irreduceable intuitive thinking down into a series of logical steps. If it is assumed that ethics is a cognitive undertaking, it has to be formalized in a series of moral rules. But one or more of these rules can take precedence over others in certain situations. In this case the expert system should be capable of making some sort of adjudicatory judgment to deal with these rules in situations where they might conflict (LaChat, 1986). Therefore, fuzzy moral rules need not be set up in a strict hierarchy and, consequently, they cannot be represented appropriately in expert systems.

Newman (1988) also argues that when a patient visits a physician with a complaint, the decision-making process which results in the physician's prescription of a course of action for the patient stems from the training, experience and judgment of the physician, both as a professional and as a human being. Could such training, experience, and judgment be built into computers as an expert system?

The unresolved issues of liability in decision-making when using expert systems raises even more ethical as well as legal questions in case of negligence or incorrect decision making. Professions — i.e., accountants, physicians, lawyers, teachers, and scientists — are exclusive repositories and disseminators of specialist knowledge. Most professions are able by regulation to control the entry to the profession and to practice. They are usually permitted to set up their own standards and to discipline their own members. Professionals are expected to stand by the judgment they make and suffer the penalties if negligence is subsequently shown. This is a crucial element, however, when making decisions to use a particular expert system. In such situations, it is worth questioning the degree to which that expert system should be given various responsibilities.

Researchers (i.e., McCarthy and Perrolle, 1990; Frank, 1988; Sheil, 1987; Waldrop, 1987) pose the question of liability when using expert systems in making decisions. They pose, for example, the question of a physician's liability for misdiagnosis of a patient after using a medical expert system. Who is responsible: the doctor, the expert system, the expert system developer, or the experts themselves? While the general view in the medical literature is that responsibility remains vested in the physician (Newman, 1988), there is no clear answer to that issue when using expert systems in business decisions.

To further illustrate this issue, assume that an expert system is being used to serve as a bank loan examiner. Each applicant sits at a terminal in the bank's offices, answering questions about his or her financial status, while the computer verifies everything by automatic queries through the network to other banks and credit companies. Finally, the system makes a decision: yes, the applicant qualifies, or no, the applicant doesn't qualify.

This expert system application could be an efficient and useful system and certainly it would be consistent in applying the bank's loan policy to each applicant. People, however, may not be willing to put up with that kind of treatment from a computer-based system, perhaps because the situation obscures the fact that the machine's "decision" actually embodies a policy made by humans.

Sometimes a straightforward application of expert systems can easily obscure the lines of responsibility. What would happen, for example, if the decision-maker did not follow the expert system's advice or conversely, if (s)he followed the expert system's advice and it turned out to be in error (Liebowitz, 1987). It is also conceivable that some managers will attempt to absolve their actions by blaming it on the machine (McCarthy and Perrolle, 1990; Waldrop, 1987, p. 242).

Finally, it can be concluded that expert systems have no consciousness and intelligence – as a human characteristic – never exists without consciousness and self-awareness (Withington, 1987; Lloyd, 1985). Expert systems lack these attributes and, consequently, artificial intelligence must always be fundamentally different from human intelligence. Thus, the concept of artificial ethics is still science fiction.

In-biasedness

Waldrop (1987, p. 248) argues that intelligent machines will embody values, assumptions, and purposes. Therefore, it becomes imperative to think carefully and explicitly about what those built-in values are before using expert systems in a decision-making capacity.

A bank loan advisor expert system, for instance, would not be prejudiced against any specific client groups unless its data and inferential rules were biased in the relevant way. A program could be written so as to embody its programmer's prejudices, but the program can be printed out and examined, whereas social attitudes cannot (Boden, 1990). Simon (1990) adds that if some of the discriminant functions built to make the credit decision can be proved to be prejudicial to certain groups, then court cases can be made. As expert systems get into more problematic decision areas – medical diagnosis, cancer treatment – people won't be happy with a program unless they are satisfied with the balance of values that is implemented in the system.

Since the knowledge engineer is the developer of the system based on expertise supplied by the domain expert, it could be argued that the knowledge engineer could, through misunderstanding or misrepresentation, invalidate the system. Consequently, the knowledge engineer's feelings and personal opinions could find their way into the system and contaminate the knowledge base (i.e., Colby, 1986). Thus, any biases held by the knowledge engineer could also influence the way the decisions are made (Mykytyn *et al.*, 1990).

It seems that expert systems designers could somehow shape what the users do, what they see, what they know about and, consequently, influence their decisions in certain ways. In other words, the tacit assumptions and values of well-intensioned systems designers could cause decision makers to drift in a direction they might not have taken by choice.

In the computer-based nuclear war decisions, Waldrop (1987, p. 247) explains, the real choices that are available to the human decision maker will have already been made by those who prepare the items in the menu. In this case, it becomes critically important to know what their assumptions are. If they only present options that refer to this or that degree of belligerence, with no options that allow for backing away from hostilities, then the system has built into it the presumption that there will be war. It is also possible that even when expert systems are used in an advisory capacity, they could impact the decision maker's final decisions in dealing with ethical issues.

In addition, expert systems may be wrongly viewed as computer programs which don't make mistakes. Software testing and validation techniques should decrease the likelihood of errors in the reasoning used and quality of decisions reached. However, since an expert system's knowledge consists of facts and heuristics derived from the expert, it is possible that the special cases may have been overlooked when constructing the expert system. In other words, a complete set of heuristics on a particular subject may not have been acquired from the expert and encoded into the expert system (Liebowitz, 1987).

Obviously, the user of a particular expert system is very much limited to operating within the boundaries of the system's domain and the system's capability for providing meaningful explanations. These explanations would be based on the facts, judgments, intuition, and experiences of the experts and knowledge engineer's representation of the information. Withington (1987) adds that expert systems are asked to make inferences based on whatever knowledge they happen to have; not necessarily complete, correct, or up-to-date. They will, therefore, make errors. Expert systems could be perceived as quite effective systems for one user but less so, or even erroneous, for another user depending on the user's knowledge, and responses to the system's questions. Different queries and responses could cause the system to reach incorrect conclusions and make faulty recommendations (Mykytyn *et al.*, 1990).

Conclusion

Expert systems — as knowledge-based information systems — attempt to achieve expert-level results in problem solving. They are expected to have human attributes in order to replicate human capacity in ethical decision making. However, most business decisions are based on imprecise information, which limits the extent to which they can be supported through automation. Therefore, the concept of artificial ethics is still science fiction.

A computer program is a symbolic representation of the world rather than a part of the world itself and, consequently, it is in principle open to question. An expert system functions in light of its information, its inferential rules, and its decision criteria, each of which may be problematic (Boden, 1990). The lack of "human" intelligence, emotions, and values, as well as the possible existence of intentional or accidental bias cause some ethical concerns when using these systems in a decision making capacity.

Since expert systems with artificial ethics are not a fact of life, the responsibility for decision-making should not be completely abdicated to currently-used expert systems. Obviously, expert systems do best when there are specific pieces of knowledge for which they can act as a big recognition memory. These systems, however, are still not good at reasoning about fuzzy things such as moral rules.

Given the heterogenuity of the decision situations, specialized expert systems should find their natural use as assistants rather than as primary decision makers. Managers should not abandon their responsibility for evaluating and, if necessary, rejecting the advice or conclusion of expert systems. This approach casts expert systems technology in a supportive and subordinate role, which matches well the current technology's limited capabilities. Expert systems can take over routine and well-structured decisions and free the manager for more demanding problem solving.

Managers should not absolve themselves of legal and ethical responsibility on the basis that a possibly inaccurate advice is the product of an expert system. They have the obligation to examine carefully the rationale and validity of their expert systems. Most managers, however, may not be able to tell when an expert system is giving wrong advice that appears sound, but which will lead to long-term problems.

One way to overcome such a problem is to search for and specify software standards that would have to be met in developing expert systems, which would be subject to regular inspections. Software should be correct, reliable, and trustworthy, and expert systems' software can only be trustworthy when they satisfy particular moral requirements based upon their area of speciality. In addition, certain types of software could perhaps be signed off by licensed software engineers, and, consequently, they would be held accountable for their performance.

It is also important for systems designers to reconsider the balance between high-tech and high-touch in the development of expert systems. Managers may not want machines that make decisions, but rather a machine that enables them to make more accurate, reliable analysis, to obtain relevant information, to gather consultative support more rapidly, and record it more accurately. Therefore, the user interface component of computer-based systems should be emphasized as the guiding image in system design. User interface puts the emphasis on the user, not the machine (Shneiderman, 1991). Following this guiding image, the division of labor between expert systems and expert humans can be achieved in a synergistic way.

References

Andrews, K. R.: 1989 (Sept.—Oct.), 'Ethics in Practice', *HBR*, pp. 99—100.

Ansari, A. and B. Modarress: 1990, 'Commercial Use of Expert Systems in the U.S.', *Journal of Systems Management* **41**, pp. 10—13, 32.

Asimov, I.: 1964, *The Rest of the Robots* (Doubleday, NY).
Barnett, G.: 1982, 'The Computer and Clinical Judgment', *New England Journal of Medicine* **307**, pp. 494–495.
Blois, M. S.: 1980, 'Clinical Judgment and Computers', *New England Journal of Medicine* **303**, pp. 192–197.
Boden, M.: 1990, 'The Social Impact of Artificial Intelligence', in R. Kurzweil (ed.) *The Age of Intelligent Machines* (MIT Press, Cambridge, MA), pp. 450–453.
Bommer, M., C. Gratto, J. Gravander and M. Tuttle: 1987, 'A Behavioral Model of Ethical and Unethical Decision Making', *Journal of Business Ethics* **6**, pp. 265–280.
Cadbury, A.: 1987 (Sept.–Oct.), 'Ethical Managers Make Their Own Rules', *HBR*, pp. 69–73.
Churchland, P. and P. Cherchland: 1990 (Jan.), 'Could a Machine Think?', *Scientific American*, pp. 32–37.
Clark, R.: 1988 (Dec.), 'Economic, Legal, and Social Implications of Information Technology', *MIS Quarterly*, pp. 517–520.
Colby, K.: 1986, 'Reloading a Human Memory: A New Ethical Question for Artificial Intelligence Technology', *The AI Magazine* **6**(4), pp. 63–64.
Cox, E. and M. Goetz: 1991 (Mar. 11), 'Fuzzy Logic Clarified', *Computerworld*, pp. 69–70.
Dejoie, R., G. Fowler and D. Paradice: 1991, *Ethical Issues in Information Systems* (Boyd & Fraser Publishing Company, Boston).
Dreyfus H. and S. Dreyfus: 1986, *Mind over Machine* (The Free Press, N.Y.).
Ferrell, O. C. and J. Fraedrich: 1991, *Business Ethics: Ethical Decision Making and Cases* (Houghton Mifflin Company, Boston).
Firth, R.: 1952 (Mar.), 'Ethical Absolutism and the Ideal Observer', *Philosophy and Phenomenological Research*, pp. 317–345.
Fleming, J. E.: 1985, 'A Suggested Approach to Linking Decision Styles with Business Ethics', *Journal of Business Ethics* **4**, pp. 137–144.
Frank, S. J.: 1988, 'What AI Practitioners Should Know about the LAW: Part 1', *AI Magazine* **9**, pp. 63–75.
Gould, P.: 1985, 'Thinks That Machine', *Integer Psychiatry* **3**, pp. 229–232.
Jastrow, R.: 1987, 'Toward an Intelligence Beyond Man's', in H. J. Watson, A. B. Carroll, and R. J. Mann (eds.), *Information Systems for Management: A Book of Readings* (Business Publications, Inc., Plano), pp. 512–514.
LaChat, M. R.: 1986, 'Artificial Intelligence and Ethics: An Exercise in the Moral Imagination', *The AI Magazine* **7**, pp. 70–79.
Leonard-Barton, D. and J. Sviokla: 1988 (Mar.–Apr.), 'Putting Expert Systems to Work', *HBR*, pp. 91–98.
Liebowitz, J.: 1987, 'Common Fallacies Expert Systems', *Computer & Society* **16**, pp. 28–33.
Lloyd, D.: 1985, 'Frankenstein's Children: Artificial Intelligence and Human Value', *Metaphilosophy* **16**, pp. 307–318.
Mason, R. O.: 1986, 'Four Ethical Issues of the Information Age', *MIS Quarterly* **10**, pp. 5–12.
McCarthy, R. L. and J. A. Perrolle: 1990 (Sept. 13–16), 'Ethics and the Computerization of Pharmacy', *Proceedings of the Conference on Computers and the Quality of Life* (The George Washington University, Washington, DC), pp. 157–163.
McCorduck, P.: 1979, *Machines Who Think* (W. H. Freeman, San Francisco).
Mykytyn, K., P. Mykytyn and C. Slinkman: 1990 (Mar.), 'Expert Systems: A Question of Liability?', *MIS Quarterly*, pp. 27–42.
Negroponte, N.: 1990 (Sept.), 'Why Have We Failed to Come Up With "Real" AI?', *BYTE*, pp. 351–354.
Newman, M.: 1988, 'Professionals and Expert Systems: A Meeting of Minds?', *Computer & Society* **18**, pp. 14–27.
O'Brien, J. A.: 1990, *Management Information Systems: A Managerial End User Perspective* (IRWIN, Boston).
Paradice, D. B. and R. M. Dejoie: 1991, 'The Ethical Decision-Making Processes of Information Systems Workers', *Journal of Business Ethics* **10**, pp. 57–77.
Parker, D. B.: 1979, *Ethical Conflicts in Computer Science and Technology* (AFIPS Press).
Rosenfield, R.: 1991 (Dec. 23 and 30), 'Minds and Machines', *The New Republic*, pp. 40–44.
Searle, J.: 1990 (Jan.), 'Is the Brain's Mind a Computer Program?', *Scientific American*, pp. 26–31.
Sheridan, J. H.: 1987, 'Is There a Computer Criminal Working for You?', in H. J. Watson, A. B. Carroll and R. J. Mann (eds.), *Information Systems for Management: A Book of Readings* (Business Publications, Inc., Plano), pp. 499–511.
Sheil, B.: 1987 (Jul.–Aug.), 'Thinking About Artificial Intelligence', *HBR*, pp. 91–97.
Shim, J. P. and G. S. Taylor: 1988, 'Business Faculty Members' Perceptions of Unauthorized Software Copying', *OR/MS Today* **15**, pp. 30–31.
Shneiderman, B.: 1990 (Sept. 13–16), 'Human Values and the Future of Technology: A Declaration of Empowerment', *Proceedings of the Conference on Computers and the Quality of Life* (The George Washington University, Washington, DC), pp. 1–6.
Simon, H. A.: 1990, 'Information Technologies and Organizations', *The Accounting Review* **65**, pp. 658–667.
Taylor, G. S. and J. S. Davis: 1989, 'Individual Privacy and Computer-Based Human Resource Information Systems', *J. of Business Ethics* **8**, pp. 569–576.
Taylor, P. W.: 1975, *Principles of Ethics: An Introduction to Ethics* 2nd ed. (Dickenson, Encino, Calif.).
Waldrop, M. M.: 1988 (Jul.), 'Soar: A Unified Theory of Cognition?', *Science*, pp. 296–298.

Waldrop, M. M.: 1987, *Man-Made Minds: The Promise of Artificial Intelligence* (Walker and Company).

Waldrop, M. M.: 1987 (Sept.), 'Causality, Structure, and Common Sense', *Science*, pp. 1297–1299.

Walter, D. O.: 1988, 'Creativity and Expert Systems: A Divergence of the Minds', *Computer and Society* **18**, pp. 5–10.

Weizenbaum, J.: 1976, *Computer Power and Human Reason* (MIT Press, Cambridge, MA.).

Withington, F.: 1987 (Mar.), 'The Mature Intelligent Computer', *MIS Quarterly*, pp. 1–3.

University of Massachussets,
College of Business and Industry,
Dartmouth,
North Dartmouth, MA 02747,
U.S.A.