A Research Perspective: Artificial Intelligence, Management and Organizations

Peter Duchessi
University at Albany, Albany, NY, USA
Robert O'Keefe
Rensselaer Polytechnic Institute, Troy, NY, USA
Daniel O'Leary
University of Southern California, Los Angeles, CA, USA

ABSTRACT **In recent years many companies have deployed Artificial Intelligence (AI), which has included neural networks, expert systems and voice-recognition systems. Yet managers and developers understand very little about how management and organizations affect or are affected by the technology. Using specific examples from practice and research, this paper discusses the interaction of AI, management and organizations, and describes some methodological approaches and theoretical models for studying those interactions. It provides direction for future research.**

INTRODUCTION

Artificial Intelligence (AI) has moved from research laboratories into business. Recent surveys indicate that a large number of companies have developed AI applications in the last two years and the growth of applications continues today (Kornel, 1990; Francett, 1991). Many of these applications are stand-alone systems, but others are integrated with more traditional Information Systems (IS), such as data-processing and Management Information Systems. Most applications are knowledge-based Expert Systems (ES), but there is a growing number of applications of other AI technologies, such as neural networks, knowledge-based planning and scheduling systems, speech-synthesis systems and voice-recognition systems (Feigenbaum *et al.*, 1988; Andrews, 1989; *Business Week*, 1992; Murphy and Brown, 1992).

Despite the proliferation of the technology, managers and developers understand little about the practical issues associated with the interaction of AI, management and organizations. This is an important topic because the success of an AI system depends on the resolution of a variety of technical, managerial and organizational issues; yet academic research is limited. O'Leary and Turban (1987) examined theoretical foundations for assessing the impact of AI on organizations. Some researchers discussed the organizational impact of ES by performing an analysis of a single system (e.g. Sviokla, 1990) or by comparing a group of systems (e.g. O'Keefe *et al.*, 1993). Others have analyzed the ES implementation process to develop an understanding of 'critical success factors' and to provide managers with guidelines for achieving a successful implementation (Irgon *et al.*, 1990; Meyer and Curley, 1991; Duchessi and O'Keefe, 1992).

Studies of the interaction of AI, management

1055-615X/93/030151-09$09.50
© 1993 by John Wiley & Sons, Ltd.
Received 24 November 1992
Revised 1 March 1993

and organizations are less numerous than with other IS, for at least two reasons. First, relative to traditional IS, most AI applications are new (mostly in the last 5–10 years), limiting our collective knowledge about their organizational impact. Second, in the early to mid-1980s technical development issues (e.g. knowledge acquisition, programming and system validation) predominated over the broader, non-technical organizational issues. Thus, the interaction of AI, management and organizations is a young field of inquiry. The purpose of this paper is to discuss the nature of the interaction, explore the design of pertinent research and suggest some related research opportunities.

The next section presents a general framework for discussing the interaction of AI, management and organizations; we use specific examples from practice and research to illustrate portions of the framework. We then summarize the prominent methodological approaches and theoretical models for analyzing AI, management and organizations, respectively. Finally, we provide direction for future research.

AI, MANAGEMENT AND ORGANIZATIONS

One framework for discussing the interaction of AI, management and organizations appears in Figure 1. Although organizations and management obviously interact directly (Orlikwoski, 1992), we emphasize the issues directly associated with AI.

Organizations are characterized by their institutional properties, including structure, size and performance. These factors provide different contexts for the development and implementation of AI, and reflect the positive or negative consequences of the technology. Management plays a critical role in admitting and supporting the technology (e.g. through provision of resources), and may use it as a business strategy. AI is not the same product in all situations, it shapes and is shaped by the framework's other two components.

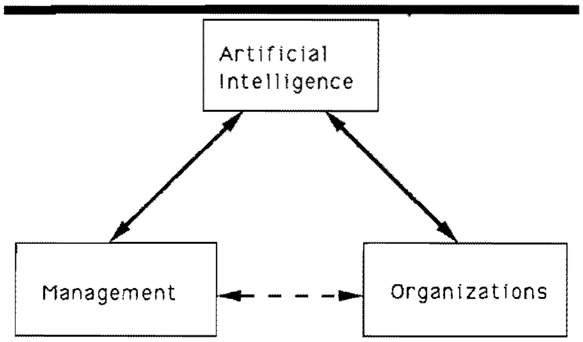

Figure 1 A simple framework for considering the interaction between AI, management and organizations.

The Impact of AI on Organizations

Some of the effects of AI on organizations include: power shifts; reassignment of decision making responsibility; cost reduction and enhanced service; and personnel shifts and downsizing. Here we review these conspicuous effects, recognizing that there are many others.

Power Shifts

The possibility of power shifts within an organization due to change in the ownership and control of knowledge has been discussed (O'Leary and Turban, 1987). As an example, PC Call Screener, developed by Eastman Kodak in the late 1980s, is an ES that diagnosed common problems with personal computers, including display, disk drive and communication problems. It allowed clerical personnel to assist users over the phone, eliminating the need for some on-site service calls by technical specialists. Implementation of the system revealed that clerks with the system solved more problems than technicians without it, and that technicians engaged in unnecessary tangential thinking. The system gave clerks the ability to assume the roles of more highly skilled technicians, reducing the power of the later group (Duchessi and O'Keefe, 1993).

Reassignment of Decision-making Responsibility

AI has the ability to change the ownership and responsibility for decision making. An example of this is American Express's Authorizers Assistant, an ES that handles the vast majority of requests for expenditure authorization made with the American Express card. The system allowed American Express to automate much of its credit authorization responsibility, removing the ownership of the decision from human authorization clerks. In the area of personal loan and credit analysis, neural networks are now being used by many major credit card companies, including Citibank and General Electric Financial Services, to perform some of the credit-granting decision making. Corporate secrecy means that details about these systems and their use are scarce.

Cost Reduction and Enhanced Service

Implementation of AI systems can help reduce costs, enhance a service provided by the organization, or do both. In addition to automating authorization decision making, the Authorizers Assistant has allowed American Express to greatly reduce labor costs and better manage its provision of a card with no fixed limits. These types of business benefits are now more acclaimed by management than the conventional benefits (including reduced decision making time, better use of expert time and codification of knowledge).

Personnel Shifts and Downsizing

AI can contribute to an organization's software maintenance expense and often requires a dedicated support staff. Although this is the case for other IS, given the dynamic nature of knowledge the cost of maintenance and enhancement of AI applications may exceed that of traditional IS. It was rumored that XCON, at one time, had a full-time staff of 50 dedicated to its maintenance. Regarding major downsizing, in 1992 AT&T announced that it would replace up to one-third of its 18 000 operators when a new AI-based speech-recognition system was installed (*Wall Street Journal*, 4 March 1992). This is the first example of major job losses due to the implementation of AI. These examples demonstrate that AI can increase the number of overhead employees and reduce the amount of direct labor, and will typically result in both occurring.

The Impact of Organizations on AI

Organizational characteristics, including job design, process design and culture, affect the deployment of AI systems. For example, O'Leary and Watkins (1992) indicate that certain organizational characteristics (e.g. size, technological awareness and IS budget) influence adoption of ES. The same system may be implemented and/or used differently in the context of different structures and cultures. We consider the issue from several viewpoints: user incentives to adopt AI; external organizations; organizational structure; and organizational support. We recognize that we are only scratching the surface of a complex issue (perhaps even less well understood than the impact of AI on organizations).

User Incentives to Adopt AI

As with other types of IS, AI systems are unlikely to be used if users do not have an incentive to adopt them. CLASS (Commercial Loan Analysis Support System) provides commercial loan officers with support to evaluate a company's financial position, recommend loan convenants and document the commercial loan analysis. Although the system demonstrated technical expertise, it provided few incentives for loan officers to use the system. CLASS required loan officers to use computers in their problem solving and this did not fit with the corporate culture that precluded the use of computers in a loan officer's office. Moreover, the loan officers never developed a personal stake in the system. The system confirmed their opinions, but never demonstrated personal gains for the loan officers, even though they agreed that the system would help them avoid bad loans. These factors had a significant impact on their tacit decision not to use the system (Duchessi and O'Keefe, 1992).

External Organizations

As AI systems become larger and more visible, the possibility for outside organizations (including unions and regulatory agencies) to have an impact on their development and deployment increases. The AT&T speech-recognition system, introduced above, offers a rare example of how an external organization can affect an AI implementation. AT&T announced the implementation of its speech-recognition system during contract negotiations with the operators' union, Communication Workers of America (CWA). The system became a significant factor in labor relations and negotiations. Eventually, CWA negotiated a settlement that has AT&T giving operators, who were replaced by the new system, a 'crack' at other jobs within AT&T (*Wall Street Journal*, 3 July 1992).

Organizational Structure

Drucker (1988) suggests that organizations are moving away from the classic stovepipe structure. With the emergence of self-managed teams, distributed responsibility and decentralized structures, there are new opportunities for AI because the technology facilitates decentralized decision making, more consistent decision making and greater reliability in decision processes. Mrs Fields, Inc. uses ES to help manage its network of retail stores (Pancari *et al.*, 1991). The systems are used to project Debbie Fields' (the founder) influence into the stores, allowing field managers to run the stores in the same way that she ran her first store 10 years ago.

Organizational Support

Users, their immediate management and ancillary support staff hold the power to advance or inhibit AI systems. Any one of these may reduce operational use by limiting the number of users, changing the composition of the target group, withholding resources and/or restricting the area to be affected within the organization. Duchessi and O'Keefe (1993) found that organizational support, as measured by turnaround of users' requirements, adequate computer resources and general community support, has a positive impact on operational use.

The Impact of AI on Management

Focusing on ES for the moment, ES that result in product/service differentiation and/or cost reduction have a direct impact on management strategies for gaining competitive advantage. With some ES, management can take offensive or defensive actions for coping with competitive forces and create a defensible position for the company. The strategies need not be limited to just products, but can include other actions, such as the development of a well-trained workforce.

Products

Perhaps the most prevalent examples of using AI for product differentiation are Expert Configuration Systems, which take a product specification or description and generate a parts list and instructions for putting the product together. The father of such systems, XCON, has been written about extensively. Similar examples outside the computer industry include Carrier's EXPERT system, which produces designs for large complex air-conditioning units for multi-story buildings, and General Electric's Computer-aided Requisition Engineering (CARE) system, which allows salespersons to search a database for electric motors that meet customer specifications, or automatically design a new motor if there are no existing ones. These systems result in fewer engineering errors, reduce base costs and reduce cycle time. However, the major impact is the company's ability to offer (in the words of Digital) *à la carte* manufacturing to its customers, who no longer have to choose a model with a few options because they can specify what they want and have it made just for them.

Workforce

For many service organizations, maintaining and effectively using a skilled workforce is crucial to profitability. For example, public accounting firms must disseminate changes in tax and accounting information to each of the accountants that perform those activities. Coopers and Lybrand's ExperTax system guides accountants through the information-gathering process and helps them explain differences between statutory and effective (or computed) tax rates (Shpilberg and Graham, 1989). The system notes relevant issues, describes the importance of information requested and analyzes it to identify critical issues for audit and tax managers. Willingham and Ribar (1988) consider auditor training to be a primary benefit of ES. These types of systems can reduce labor costs, increase accuracy and provide product differentiation for basic 'vanilla' services, especially when the customer has little marketing information to differentiate providers.

The Impact of Management on AI

Management plays a key role in admitting AI into the organization and implementing it successfully. The two primary ways management exerts its influence are: acting as a champion for an AI system and providing resources for its implementation.

Champion

One of the primary issues in the implementation of AI systems seems to be the need for a champion to promote the use of AI (Hayes-Roth *et al.*, 1983). Being a champion goes beyond just verbal support for the system, it includes a willingness to actively advocate the technology and make it a high priority in the organization. Duchessi and O'Keefe (1993) found that top management support and manager acceptance are important to ES implementation success, and favorably impact users' perception of management support and operational use.

Provision of Resources

O'Leary and Watkins (1992) found that organizational pressure to adopt ES, management support and provision of adequate budgets for the technology are positively related to ES adoption. Duchessi and O'Keefe (1992) discovered that management support includes provision of people, time and money. By being the initial source of support, committing resources and making the implementation at least as important as other business activities, management can have a significant positive impact on successful deployment of AI.

AI and Other Information Technologies

Although AI has been treated as an independent technology, it is a single component of the IS portfolio. IS planning attempts to align computer-based systems with the needs of the organization. Top-down approaches begin with business objectives and derive desirable architectures to support those objectives which encompass all aspects of computing, and will determine (at least in part) the nature and number of AI systems on an employee's desk. Through IS planning, the organization selects what AI systems it wants to build, establishes the level of funding for them and determines how they should be integrated with existing and planned databases, data entry systems, reporting systems, and decision support technologies. Organizations with aggressive plans that include AI will deploy the technology throughout their businesses. Ultimately, the contribution of AI systems are more likely to be a function of how they integrate and interface with other hardware, software, policies, procedures and organizational arrangements that collectively constitute an improved business process. Thus, AI systems will be components of larger business systems with confounding cost and benefit issues.

METHODOLOGICAL APPROACHES

The two predominant approaches for studying the interaction of AI, management and organizations are case studies (both single and multiple) and empirical studies. Most of the case studies are about successful applications; there is a dearth of cases about failures which may be even more important to analyze relative to successes. Cases provide both practical insights and a basis for developing theories to be eventually tested by empirical methods.

Case Studies

To date, analysis of a single case has been the most often used approach to studying AI, management and organizations. There are a number of published studies that focus on a single successful application. For instance, Sviokla (1990) described the organizational impact of XCON at Digital. He found that the use of XCON increased the information-processing capacity of the organization, the system altered the local execution of the configuration task and the system directly supported Digital's product strategy.

In contrast, Reitman and Shim (1993) used a case analysis to discuss customer and vendor viewpoints of an unsuccessful implementation of Palladian Software's Management Advisor. With regard to the customer, they found companies attempting to implement an ES for financial planning must have an adequate understanding of the strategic and/or organizational problems to be addressed by the system. Additionally, developing high-end commercial ES for strategic financial applications entailed both technical and market risks for the vendor. The more novel the system, the broader its scope, and the greater the discrepancy between the task interactions supported by the system and the client's actual processes, the greater the risks.

Cases about failures are notable because they provide a different perspective. They sometimes include factors that are also found to be present in analyses of successful AI implementations, questioning the importance of those factors as contributing to success. Moreover, they are extremely rare: managers, project leaders and/or developers are, in general, reluctant to admit to and discuss failures.

To analyze ES implementation success or failure, Duchessi and O'Keefe (1993) used a multiple case design where each case serves as a separate experiment that confirms or disconfirms inferences drawn from the group of cases. After each case is analyzed separately, the method organizes the cases into success and failure categories to facilitate a cross-case search for patterns. The search focuses on identifying within-group similarities among successful and less successful categories as well as inter-group differences. The cases are revisited to determine if they support propositions emerging from the search process. The multiple case design is useful for identify propositions or constructs for building theory. Eisenhardt (1989a) provides a good description of the method for inducting theory and has successfully employed it elsewhere (Eisenhardt, 1989b).

Empirical Studies

There are a few empirical studies on AI, management and organizations, and the existing studies focus on where and how the technology is being used in the marketplace. Generally, consulting firms and AI vendors perform the surveys to determine the extent of AI development and usage in organizations, and thus follow quite narrow research perspectives. Here we briefly summarize two typical academic studies.

Pickett and Case (1990) surveyed R&D professionals on AI/ES applications in R&D. The survey revealed that companies with the resources to deploy the technology are doing so very cautiously, and view the technology as a means to capture irreplaceable expertise and improve control over complex systems. Additionally, the respondents identified several impediments to using AI and ES, including selling management on the value of the technology, development of a knowledge base and lack of suitable development tools. The small sample size (33 responses) limit the study's value.

Doukidis and Paul (1990) performed an empirical analysis of the application of AI techniques among members of the UK Operational Research (OR) society. The study's findings include: professional and organizational motivation is the main reason for using AI; other departments within an organization approach OR departments for AI development, demonstrating the success of some OR departments at appearing to be innovative; and ES are the major AI technique employed because of their cost-effectiveness.

To date, empirical studies have been descriptive rather than inferential in nature. Few studies are conceived as well-defined research projects to test propositions and theoretical models that emerge from careful analysis of the extant literature or case studies. However, the excessive cost, time and difficulty in obtaining a representative sample are strong constraints to performing an empirical study.

THEORETICAL MODELS FOR ANALYSES

There are a number of theoretical models that can be used to investigate the interaction of AI, management and organizations. For example, AI represents a sophisticated, high-level type of IS, and thus one or more existing IS implementation models may offer useful perspectives for analyzing AI implementation. To date, there is no core set of constructs because existing models focus on a limited set of variables. Reviews of several potential models from several disciplines appear below.

Socio-technical

Based on a literature review, Sharma *et al.* (1991) proposed a socio-technical model that seeks to address some of the important questions behind ES deployment such as: What are the procedures that facilitate successful implementation? Under what circumstances do positive effects materialize? What are some fundamental causal relationships? The model simultaneously addresses the technical dimension, including task domain, computer platform and knowledge engineering process, and the social dimension, including user interaction, manager support and organizational fit. The model suggests that quality of an ES is a function of the nature of the task, applied technology, support from people, organizational parameters (e.g. culture, structure and external environment) and the associated interaction among these components.

Management Strategy and Structure

The management literature provides a number of viewpoints on the determinants of organizational structure. Chandler (1962) is responsible for the classic work on organizational strategy and structure, and heavily influences the current work in this area. According to Chandler, companies develop new structures to meet administrative needs, which result from an expansion of a company's activities into new areas, functions or product lines. A new strategy requires a new structure, or a refashioned one, to maintain or enhance organizational efficiencies.

A number of companies, including Texas Instruments, Arthur Andersen and Fujitsu, use AI strategically either to enhance the efficiency of their operations and/or to sell AI systems as new products (Feigenbaum *et al.*, 1988). In these types of companies the theory on strategy–structure relationships provides a basis for analyzing the interaction of AI, management and organizations. For example, companies that aggressively pursue AI may require structures that promote environmental scanning, flexibility and lateral communications.

Organizational Innovation

If the deployment of AI is viewed as a technical innovation, then organizational innovation models are especially appropriate for studying the interaction of AI, management and organization. According to Kwon and Zmud (1987), organizational innovation can be viewed as a three-stage process: initiation, adoption and implementation. Initiation results from pressure to change, adoption involves provision of resources and implementation refers to the development, installation and maintenance activities. Others have expanded the implementation phase to include acceptance, usage, performance, satisfaction and incorporation (e.g. Rogers, 1983; Schultz *et al.*, 1984). These models are valuable because the phases address specific technical, motivational and political issues, and incorporate a number of associated constructs.

Task-based

Since the focus of much implemented AI is on the automation or support of specific tasks, then models that focus on the task-based level of management, rather than the broader organizational context, are relevant for giving insight into the adoption and implementation of AI. This can allow for generalization across certain task domains, and perhaps even occupational groups, but as these models are below the organizational level they are not appropriate for analyzing entire organizations.

The work of Perrow (1967) is a well-known example of this approach. He argued that tasks, and perhaps even entire occupations (e.g. accountants, engineers) or parts of common occupational work (e.g. tax accounting), can be meaningfully differentiated based upon the way they use information and knowledge to perform tasks and handle exceptional instances of those tasks. O'Keefe *et al.* (1993) have used these ideas in an analysis of implemented ES in accounting, and show some differences between tax and auditing ES that would be expected from the model.

Information Systems Implementation

There is considerable research that examines the factors associated with the successful adoption, development and implementation of IS. Lucas *et al.* (1990) proposed a model of IS implementation that consists of manager and user models. The manager model includes top management support, management belief in system concept and manager–researcher involvement, while the user model includes user knowledge of system purpose, user's personal stake and user's job characteristics. The model also relates the factors to one another, and is appealing to AI and ES researchers for three reasons. First, it integrates previous findings on IS implementation research. Second, it is a view based on the relationships between factors and not simply the presence or otherwise of such factors. Third, the model is two-stage, one stage representing management initiation and support, the other representing user acceptance and use.

Information Systems Implementation and Organizational Innovation

Kwon and Zmud (1987) combine stages of the organizational innovation process (i.e. initiation, adoption and implementation) with IS implementation factors, such as individual, structural, task and environmental factors, to develop a model which provides a more complete perspective than is found in either of the organizational innovation or IS implementation models *per se*. By integrating these two streams of research, the model provides a basis for examining multiple factors associated with implementation, innovation and diffusion. This model is especially appropriate when considering AI as both an IS implementation and a technological innovation.

DIRECTIONS FOR RESEARCH

Using a simple framework that focuses on AI and its separate interaction with management and organizations, we have discussed the interaction of AI, management and organizations, and presented a number of examples to illustrate the nature of those interactions. We have briefly described the predominant methodological approaches for studying AI in organizations, namely case and empirical studies. Finally, we presented several theoretical methods in the IS implementation, organizational innovation and management literature that are pertinent for forthcoming research.

We suggest three directions for future research. First, existing theoretical models should be re-examined in the context of AI and augmented with the propositions that have emerged from existing case studies, so as to develop a more comprehensive model of AI innovation and implementation. Sharma *et al.*'s (1991) effort is a step in this direction. Second, specific critical factors should be investigated across multiple case studies and organizations, so as to better understand the impact of the presence (or absence) of the factors associated with development, implementation and adoption. The investigation should include factors that have been shown to be important in the IS literature (e.g. top management support) and those that are perhaps new to AI deployment (e.g. 'experts'). Third, both case and empirical studies to date have focused on systems. Studies that focus on an entire organization (or a sizable part of one) will give insight into how AI is admitted and deployed throughout an organization. For instance, a case history on the life cycle of an internal AI group might provide considerable insight into these issues.

Whatever the research issues investigated and the methodology used, the proliferation of AI technology is providing numerous systems to study. It appears to be a good time to advance research on the topic of AI, management and organizations.

References

Andrews, B., 'Successful expert systems', *Financial Times Management Report* London, 1989.

Business Week, 'The new rocket science', 2 November 1992, 131–40.

Chandler, A., *Strategy and Structure*, MIT Press, Cambridge, MA, 1962.

Doukidis, G.I. and Paul, R.J., 'A survey of the application of artificial intelligence techniques within the OR society', *Journal of the Operational Research Society*, **41**(5), 1990, 363–75.

Drucker, P., 'The coming of the new organization', *Harvard Business Review*, **66**, No. 1, 1988, 45–53.

Duchessi, P. and O'Keefe, R.M., 'Contrasting successful and unsuccessful expert systems', *European Journal of Operational Research*, **61**(1/2), 1992, 122–34.

Duchessi, P. and O'Keefe, R.M., 'Understanding expert system's success and failure', Working paper, Decision Sciences and Engineering Systems, Rensselaer Polytechnic Institute, Troy, NY, 1993.

Eisenhardt, K.M., 'Building theories from case study research', *Academy of Management Journal*, **14**(4), 1989a, 532–50.

Eisenhardt, K.M., 'Making fast strategic decisions in high-velocity environments', *Academy of Management Journal*, **32**(3), 1989b, 543–6.

Feigenbaum, E., McCorduck, P. and Nii, P., *The Rise* *of the Expert Company*, Times Books, New York, 1988.

Francett, B., 'AI (quietly) goes mainstream', *Computerworld*, **25**(30), 1991, 59–60.

Hayes-Roth, F., Waterman, D. and Lenat, D., *Building Expert Systems*, Addison-Wesley, Reading, MA, 1983.

Irgon, A., Zolnowski, K.J., Murray, M. and Gersho, M., 'Expert systems development: a retrospective review of five systems', *IEEE Expert* **5**(3), 1990, 25–39.

Kornel, A., 'Investing in R&D prowess', *Computerworld*, Supplement, 18 August 1990, 28–9.

Kwon, T.H. and Zmud, R.W., 'Unifying the fragmented models of information systems implementation', in Boland, R.J. and Hirscheim, R.A. (eds), *Critical Issues in Information Systems Research*, John Wiley, New York, 1987, pp. 135–56.

Lucas, H.C., Ginzberg, M.J. and Schultz, R.L., *Information Systems Implementation: Testing a Structural Model*, Academic Press, Norwood, NJ, 1990.

Meyer, M.H. and Curley, K.F., 'Putting expert systems technology to work', *Sloan Management Review*, **32**(5), 1991, 21–31.

Murphy, D. and Brown, C., 'The uses of advanced information technology in audit planning', *International Journal of Intelligent Systems in Accounting, Finance and Management*, **1**(3), 1992, 187–94.

O'Leary, D. and Turban, E., 'The organizational impact of expert systems', *Human Systems Management*, 7(1), 1987, 11–19.

O'Leary, D. and Watkins, P., 'Internal auditing and expert systems: technology adoption of an audit judgement tool', unpublished paper presented at the National Meeting of the American Accounting Association, August 1992.

O'Keefe, R., O'Leary, D., Rebne, D. and Chung, Q., 'The impact of expert systems in accounting: system characteristics, productivity and work unit effects', *International Journal of Intelligent Systems in Accounting, Finance and Management*, 1993, this issue.

Orlikowski, W., 'The duality of technology: rethinking the concept of technology in organizations', *Organizational Science*, **3**(3), 1992.

Pancari, D., Senn, A. and Smiley, G., 'Are retailers sold on expert systems?', *Chief Information Officer Journal*, 3(5), 1991, 10–14.

Perrow, C.A., 'A framework for the comparative analysis of organizations', *American Sociological Review*, **32**(2), 194–208.

Pickett, J.R. and Case, T.L., 'Expert systems and artificial intelligence', *Research-Technology Management*, **33**(3), 1990, 35–8.

Reitman, W. and Shim, S., 'Expert systems for evaluating business opportunities: implementing the management advisor at Krypton Chemical', *International Journal of Intelligent Systems in Accounting, Finance and Management*, 1993, this issue.

Rogers, E.M., *Diffusion of Innovations*, Free Press, New York, 1983.

Schultz, R.L., Ginzberg, M.J. and Lucas, H.C., 'A structural model of implementation', in Schultz, R.L. and Ginzberg, M.J., (eds), *Management Science Implementation*, JAI Press, Greenwich, CT, 1984.

Sharma, R.S., Conrath, D.W. and Dilts, D.M., 'A socio-technical model for deploying expert systems — Part I: The general theory', *IEEE Transactions on Engineering Management*, **38**(1), 1991, 14–23.

Shpilberg, D. and Graham, L.E., 'Developing ExperTAX: an expert system for corporate tax accrual and planning', in Vasarhelyi, M.A. (ed), *Artificial Intelligence in Accounting and Auditing*, Markus Weiner, New York, 1989, pp. 343–72.

Sviokla, J., 'An examination of the impact of expert systems on the firm', *MIS Quarterly*, **14**(2), 1990, 127–40.

Wall Street Journal, 'AT&T to replace as many as one-third of its operators with computer systems', 4 March 1992, A4.

Wall Street Journal, 'AT&T union pact boosts job security', 3 July 1992, C21.

Willingham, J. and Ribar, G., 'Development of an expert audit system for loan loss evaluation', *Auditor Productivity in the Year 2000*, Arthur Young, Reston, VA, 1988.