Emerging Technologies

Integration of Intelligent Systems and Conventional Systems: Requirements for Co-ordinating Multiple Agents for Diagnostic Decisions

Daniel E. O'Leary and Paul R. Watkins
School of Business, University of Southern California, Los Angeles, CA, USA

Abstract This paper describes the development of an integrated intelligent (expert) system and conventional information system to support decision making of a diagnostic decision in a multi-agent environment. It argues that many *a priori* requirements for the information system can be elicited by accounting for the specific nature of the diagnostic decision and for the single/multiple-actor environment. This finding is generalized to a framework that indicates that *a priori* requirements take into account the basic theories on which the information system supporting the decision task is ultimately based. In particular, attention is focused on the nature of the task (e.g. diagnostic), whether the application is for a single- or multiple-agent situation and the nature of the multi-agent structure (e.g. product or functional form of organization).

Introduction

In most business situations expert systems are not independent of related information systems. However, few reports of expert (intelligent) systems projects include the development of an intelligent system *and* related information systems. In most summaries the focus has been on the intelligent system. The ties to a supporting information system, if any, have generally been ignored. Frequently this is because the system has been constructed to directly solicit information from the user and/or make most recommendations directly to that specific user, to support independent decision making.

Although such architectures may be appropriate for many situations (for example, medical diagnosis) a number of business problems are based in organizations. Typically, such organizations have multiple actors involved in decision processes. Information systems can provide co-ordination of those multiple actors to support the use of the expert system. Without an information system the actors may not be aware of the actions and outcomes of actions of other actors.

Thus, the purpose of this paper is to discuss the integration of an expert system into a multiple-actor organization decision-making situation, and the resulting development of an information system to support the intelligent system and those multiple actors. It is found that the *a priori* requirements for the design of that information system are dictated by the requirements of both the generic nature of the expert system/decision process (in this case, diagnostic) and the co-ordination of a multiple-

1055-615X/92/020135-11$05.50

Received November 1990
Revised September 1991

actor environment (and the structure of that multi-agent organization, e.g. functional or product-based).

Background

Initially, the purpose of this study was to develop an expert system to assist individual diagnosticians with communications problems in a micro-to-mainframe computer environment. Although the initial project goal was to build that expert system to solve a specific problem, it rapidly became apparent that the latter was embedded in a larger information systems problem. Supporting the diagnosis of communication problems required more than rule-based decision-making assistance. Co-ordinating case analyses of current problems with previous analyses of similar problems was required (case-based reasoning). Further, co-ordinating the actions of the diagnosticians with respect to a given client or problem was also necessary.

Thus, an information system was needed to co-ordinate and support the users of the expert system. The information system played an integral part in the retrieval of information about different clients, the analysts who had dealt with the clients, the problems those clients had in the past, and the solutions generated by those analysts for those past clients. The system had to serve the role of summarizing previous cases and co-ordinating the solution of current ones. In addition, the system had to function as a disperser of new information regarding newly recognized problems, such as software bugs.

Field Study Research

This research employs a field study methodology, which has been used in previous research into intelligent systems. For example, Watkins and O'Leary (1987) employed a field study of the development of expert systems by experts.

As discussed by Kaplan (1964), field research can provide important insights:

> Much of . . . science is occupied with field studies, that is, with direct or indirect observation of behavior in the circumstances in which it occurs without any significant intervention on the part of the observer. . . . Of course in his role as a consulting expert the . . . scientist may be taking part in some part of an experiment. We should not overlook the opportunities for genuine and important scientific work afforded by this social role.

Structure of the Paper

This paper is organized as follows. The next section provides a short summary of some related research in requirements analysis. The third section investigates the particular diagnosis-based task addressed in this paper and the specific information and knowledge-based systems requirements. The fourth section briefly analyzes the integration of intelligent and conventional information systems. In addition, it presents the resulting design of the integrated system on which the field study is based. The fifth section analyzes the general nature of diagnosis-based decisions and systems, and the generic information requirements for diagnosis-based decisions. The sixth section investigates the information requirements associated with the multiple-agent environment in which the system has been developed. The seventh section extends the discussion of the fifth section to other types of decisions, rather than just diagnostic, with particular focus on monitoring and control systems. The eighth section extends the discussion in the sixth section to include the structure of the multi-agent environment. The ninth section integrates the discussions in the seventh and eighth sections to provide a framework for the *a priori* analysis of requirements for information systems supporting intelligent systems. The final section provides a brief summary and some extensions.

Requirements Analysis

Requirements analysis is the process of determining what is required of an information system to assist in the design and development of that system. There have been a number of approaches suggested and employed in the development of requirements analysis for information systems. In a recent paper, Wetherbe (1991) summarized a number of those techniques (interview, analysis, prototyping and other processes).

Requirements analysis typically proceeds with an iterative investigation of the needs of the users of the information system (e.g. critical success factors) to establish the information system requirements. However, it is the contention of this paper that certain *a priori* characteristics of the specific decision task and environment must also be met. Those characteristics form information system requirements.

A Priori Requirements

System requirements can be established based on the basic nature of the system being developed. The nature of the task investigated (e.g. diagnostic) has a predictable set of information requirements that must be met. Similarly, the environmental setting in which the intelligent system is being used (single- or multi-agent) has a predictable set of information requirements that must be met. These are *a priori* in that they can be known to the designer prior to an iterative requirements analysis process.

Requirements Analysis and Knowledge Acquisition

The determination of requirements for intelligent systems has generally fallen under the topic of 'knowledge acquisition', which is concerned with the elicitation and representation of knowledge, not the information to support that knowledge for the intelligent system. Requirements analysis is used for developing the supporting information system. The differentiation between knowledge acquisition and information requirements analysis is made because of the concern with the determination of 'knowledge' necessary to build the intelligent system, in contrast to the 'information' required to build the information system to support an intelligent system.

A Diagnostic Task and its Information Requirements

The specific field study was carried out at a firm that provides tax compliance and planning software to their clients, primarily in the form of micro-to-mainframe communications. Virtually all the firm's clients have communication capabilities (micros with modems) in order to communicate with the firm's computer facilities.

As a result, the company employs a base of technicians (diagnosticians) to assist in the communication and use of the software. The technicians are centrally located and users refer all questions to that location. All questions are answered over the phone, and either the questions are solved at the time of the call or the technician calls the client back. Sometimes this is an iterative process, with the technician offering a solution, getting feedback from the user and then generating another solution. Generally, calls are handled by the agents that are available, although there is an effort to have the agent who has dealt with the client in the past handle those specific clients.

Since the diagnosis is offsite, information for diagnosis purposes must be gathered by someone other than the diagnostician. As most of the firm's clients are not sophisticated computer users, this can complicate the information-gathering process.

Unfortunately, the technicians often come into their role without adequate background or previous training. In addition, the underlying technology often changes and there is a frequent turnover of technicians. When the technicians understand the technology, their experience becomes attractive to other departments and firms. As a result, the company must continually educate new and existing technicians. Thus, the firm decided that it was desirable to have an expert system to assist the technicians in their decision making.

Diagnosing Communications Problems

The initial purpose of the expert system was to provide assistance to diagnosticians in their analysis of communications problems (e.g. modems). This task was chosen as the specific client firm's initial effort at developing an expert system. Since the problem was a diagnostic one with a relatively bounded set of diagnoses it was seen as offering a high probability of successful development. Development and knowledge acquisition for this system was discussed in O'Leary and Watkins (1990).

In order to provide appropriate response to the client the technician has both knowledge and information needs that must be met by various information flows. In addition, the organization wishes to co-ordinate the diagnosis efforts of the individual agents.

Knowledge Needs

Specific knowledge and decision needs were captured in a rule-based expert system designed to assist decision making. First, in order to be able to diagnose problems the technician must understand the nature of the technology and its changes. Thus, the system contains knowledge about the technology (e.g. modems) and changes in it. Second, in order to diagnose communications problems, technicians need an understanding of the diagnosis process in this specific context. Thus, they must understand relationships between symptoms, causes, communication problems and solutions to those problems.

The knowledge required was obtained from a number of sources. Documentation provided knowledge regarding specific communications hardware (e.g. modems). A unique form of protocol analysis led to the determination of some of the relationships between that knowledge, symptoms of problems and problem solutions. Observation provided insight into the physical activities associated with the diagnosis process. For example, it was found that technicians need to be able to visualize problems. Accordingly, visual aids were included to help guide them. With both observation and interview, it was found that technicians need to interact with other technicians to find out who has dealt with a client previously and what problems had occurred. As a result, the multiple-agent co-ordination problems became apparent. One solution to these problems was an information system to co-ordinate those efforts.

Information about Clients

The client information needs derive in large measure from the diagnosis nature of the problem. In particular, the requirements in the diagnosis of these communications problems include the following. First, the technician needs initial state data. For example, how is the client's system configured, in terms of modem, brand of computer, etc.? Second, the technician requires information relating to changes in state variables. For example, what problems has the client had in the past? Have they recently purchased new equipment? Are they a new client? What do we know about the client contact? Are they computer literate? Third, the technician needs to know about general system problems that are occurring with use of the system that may impact the client. For example, there may be an error in the software that is causing problems with communications for all clients.

Co-ordinating Multiple Agents

In addition, since multiple technicians may interact with a client there are information needs relating to the co-ordination of those multiple actors. In particular, a history of previous client communications (cases of problems and recommended solutions) and ensuring consistency between the history and the current problems are required. In addition, in the situation where no technicians are available at the time of the client's call then it is critical to ensure that at least one (and no more than one) technician follows up with the client. Otherwise, client needs will not be met or clients may feel pestered.

Integrating Intelligent Systems and Conventional Information Systems

The issue of integrating intelligent systems and other types of systems has received limited attention. Turban and Watkins (1986) discuss two perspectives of integrating expert systems and decision-support systems: expert systems as intelligent decision-support systems and expert systems as components in those systems. This paper generalizes that discussion, broadening the integration from expert and decision-support systems to expert/intelligent and information systems. In addition, this section summarizes the actual system built, using one of those approaches, to meet the requirements identified in the previous section.

Two Types of Integration

There are at least two basic approaches to the integration of intelligent and conventional systems. First, an intelligent system can be embedded in a conventional information system to support some activity in that system. For example, the security of the information system may be monitored by an expert system (O'Leary, 1990). Second, a conventional type of information system can be developed to support the information needs of the intelligent system. The focus of this paper is on the second type of integration.

System Implementation

The integrated expert system and information system is summarized in Figure 1. The initial state information is contained in the 'database of configuration'. That information is introduced into the system through the 'initializer' and updated through the 'log'. The set of history information is contained in 'customer history' and the 'log' provides a vehicle for updating customer history. The 'visual library' contains a set of diagrams (templates) of different modems. This allows the reasoning to be rapidly adapted to meet the needs of new technology, such as new modems.

The integrated system involves two approaches that are employed in intelligent systems: rule- and case-based. The rules used in this system are generic 'if *a* and *b*' ones. The case-based approach used is nonadaptive. Cases are accumulated and users can change the cases to meet their individual needs. However, there is no adaptation of cases to meet specific case needs by the system.

Generic Diagnosis Systems as a Source of Requirements

The premise of this and the next section is that requirements for the integrated system given in Figure 1 could be elicited, *a priori*, by taking

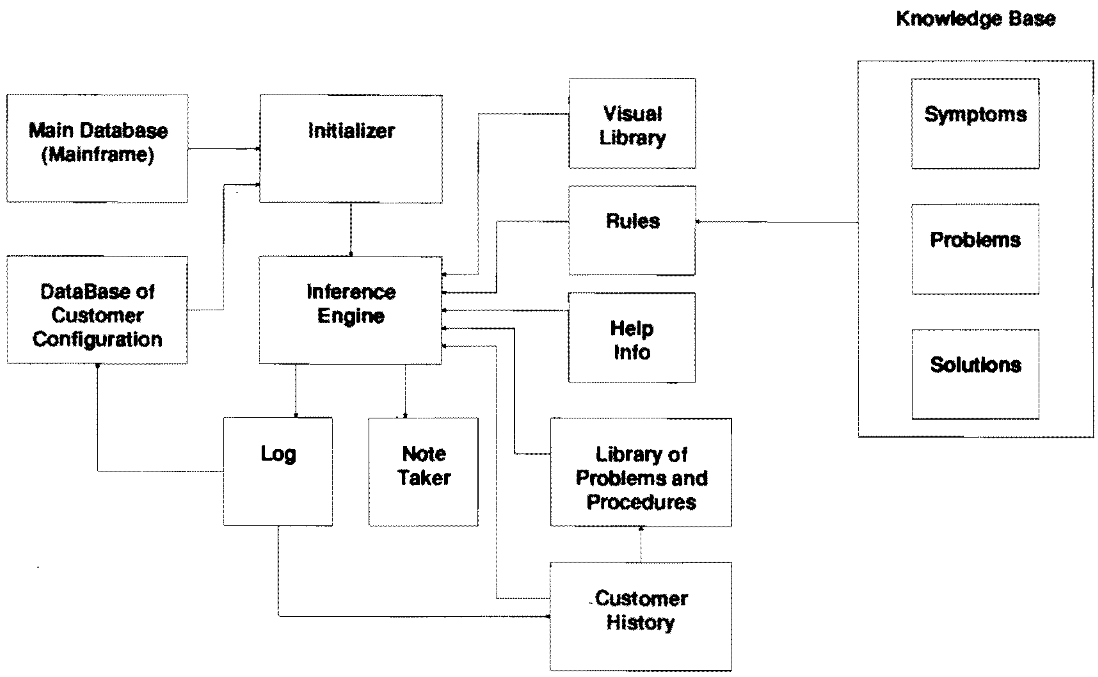

Figure 1 System architecture.

into account the generic properties of the decision problem and the multi-agent environment. In particular, this section argues that the diagnosis nature of the problem forces the need for specific types of information identified in the diagnostic application discussed earlier.

Diagnosis Problems

Diagnosis problems constitute situations where an investigator is faced with determining the cause (and often solution/repair) of a given situation. In these problems different solutions are associated with alternative patterns of information (symptoms). The quality of the diagnosis is generally regarded as a function of the ability of the diagnoser to match those symptom patterns with appropriate diagnoses.

Systems designed to solve diagnostic problems have been among the most successful applications of intelligent systems. Expert systems have been designed for business (Meservy, 1985), medicine (e.g. Buchanan and Shortliffe, 1985), and mining (e.g. Duda *et al.*, 1979). As noted by Hayes-Roth *et al.* (1983), diagnostic systems operate in a manner similar to human diagnosers: 'Diagnostic systems infer system malfunctions from observables. . . . Diagnosis systems typically relate observed behavioral irregularities with underlying causes.'

Information Requirements of Diagnosis Systems

The generic nature of diagnostic systems and decisions allows the investigation of the basic information requirements of such systems. Since diagnostic decisions are based on elicited data patterns, the diagnoser must have access to substantial amounts of data, and such data include the following. First, if the diagnosis is going to consider the basic background of the source being diagnosed then initial state data are often necessary. Second, unless each diagnosis is independent, a history of the diagnosis of previous problems may be required for the diagnostic or at least to provide insight. Third, in some diagnostic situations the basic nature of the diagnosis does not change but actual situations do. That is, the template is the same for each situation, but the content in the templates is periodically changed. Thus, the basic templates are required.

Implementation of Information Requirements

A comparison of the information requirements for generic diagnosis system reveals that the requirements for the communications system design could have been predicted. Although this is a single field study, it does suggest that critical requirements for the information systems that support diagnostic systems can be predicted from the generic decision-making nature of the system.

Multiple-Agent Co-ordination as a Source of Requirements

The premise of this section is that many of the requirements for the integrated system given in Figure 1 could be elicited, *a priori*, by taking into account the generic properties of the multiple-agent co-ordination. In particular, the multiple-agent aspect of the problem forces the need for specific types of information.

Knowledge of Agents

Bond and Gasser (1988) delineate six types of knowledge of agents that can be represented: capabilities; resources and demands; responsibilities; solution progress; communication skills; and beliefs, goals, plans and actions. These different sets of knowledge also generate information system requirements.

Information Requirements of Multiple-agent Systems

The knowledge of agent capabilities suggests capturing, evaluating and updating the performance of the individual technicians. This would allow dynamic knowledge of the current performance level of the agents interfacing with the information system.

The knowledge of agent resources and demands can be critical in scheduling the different diagnosticians to respond to client inquiries. The information system could delegate the call to an agent that is not busy. Alternatively, if an agent is busy the system could assist in the allocation of parts of the work load to other agents.

Knowledge of responsibilities is necessary to know what information should be gathered and monitored regarding which agent. Knowledge of solution progress requires that the information system monitor the progress of each agent in satisfying client requests. Such information is necessary for a number of purposes, including work load adjustment.

Knowledge of communication skills refers, in part, to knowing what information it will be useful to communicate. This is critical to performance evaluation and analyses of other behaviors of the individual agents. The knowledge of beliefs, goals, plans and actions requires the generation of a history of what each agent has done in response to various requests for action (implemented in the integrated information system as a history file).

Implementation of Information Requirements

The implementation of the information requirements to accommodate the multiple-agent nature of the decision environment exceeds the requirements that were actually built into the system. The system's primary vehicle for coordination is the use of a history file, which provides evidence that some requirements for an information system can be predicted from the theory of the environment in which it is based.

Extension: Information Systems for Other Decision Problems

Various researchers (e.g. Hayes-Roth *et al.*, 1983) have presented theoretical discussions that relate to the nature of different kinds of decisions. For example, as summarized in Table 1, those researchers suggest that a taxonomy of decision making should include diagnosis, planning, monitoring, scheduling and other types of generic decisions. Those investigations are critical to the development of expert systems, since if there are different types of decision problems and the differences in those decision types can be isolated, then it can impact expert system development processes and the information requirements.

The extent to which other information system requirements may be generated from different types of decision problems can be illustrated by examining some of those different decision types. Two specific types of systems are discussed, but other types of decisions could also be investigated.

Monitoring Systems

Consider the case of monitoring systems. As noted by Hayes-Roth *et al.* (1982, p. 15), 'Monitoring systems compare observations of system behavior to features that seem crucial to successful plan outcomes'. This indicates that the information system should provide observations of system behavior and plan outcomes and critical success features. It also suggests

Table 1 Generic categories of knowledge engineering applications.

Category	Problem addressed
Interpretation	Inferring situation description from sensor data
Prediction	Inferring likely consequences of given situations
Diagnosis	Inferring system malfunctions from observables
Design	Configuring objects under constraints
Planning	Designing actions
Monitoring	Comparing observations to plan vulnerabilities
Debugging	Prescribing remedies for malfunctions
Repair	Executing a plan to administer a prescribed remedy
Instruction	Diagnosing, debugging and repairing student behavior
Control	Interpreting, predicting, repairing and monitoring system behavior

Source: Hayes-Roth *et al.* (1983, p. 14).

that the use of a so-called 'critical success' factors approach (e.g. Rockart, 1979) may be a function of generic problem type.

Hayes-Roth *et al.* (1983, p. 15) further note that:

> Monitoring systems identify vulnerabilities two ways. One vulnerability corresponds to an assumed condition whose violation would nullify the plan's rationale. Another kind of violation arises when some potential effect of the plan violates a planning constraint.

The potential existence of violations indicates that the information system should identify any conditions that would potentially nullify the plan's rationale. In addition, the information system would need to provide constant information regarding the proximity of planning constraints to violation.

Control Systems

Hayes-Roth *et al.* (1983, p. 15) describe the functioning of control systems: 'The control system must repeatedly interpret the current situation, predict the future, diagnose the causes of anticipated problems, formulate a remedial plan, and monitor its execution to ensure success.' Such a system can make substantial demands on an information system. Interpreting the current situation requires receiving data about the current situation. Predicting the future typically involves taking data from the past and using them to model the expected future. Diagnosing current problems requires the information system to provide the information discussed immediately above. Finally, monitoring a plan needs constant information flow about the successes and failures of the implementation of the plan.

Extension: Multi-agent Structures as a Basis of Requirements

Until this point in the paper multi-agent structures have been relatively unspecified. The particular system studied as part of the field study had little multiple-agent co-ordination structure.

However, there is important research that has examined some specific multi-agent structures (e.g. Williamson, 1970; Malone, 1987). In particular, Malone employed function form, product form, centralized market and decentralized market, and argues that each of these forms and markets have different costs associated with them, i.e. production, co-ordination and vulnerability. These different structures for multi-agent systems also have varying information requirements.

Product Form: Information Requirement Issues

The product hierarchy organizes multiple-agent activity by product divisions, e.g. General Motors is divided into Buick, Cadillac, Chevrolet, Pontiac and others. In a classic product organization hierarchy there is no need for information flows between different divisions. A production success or failure in one division would not necessarily mean a success or failure in another.

Further, within a division the manager in charge must supervise a number of activities. Since managing those activities requires information about them it is likely that the information used varies, based on the manager's knowledge of the activity. For those activities where the manager has only limited knowledge we are likely to find that the depth of information will differ from those over which the manager has substantial knowledge.

Functional Form: Information Requirement Issues

The functional hierarchy is organized by similarity in e.g. processing. Accountants are organized in accounting departments, manufacturing is a department, etc. Although the different departments function independently, ultimately there is accountability of each department to the entire organization. Unlike the product form, if there is a failure in any functional component the organization as a whole may fail. Thus, there needs to be information to co-ordinate a function form organization.

In addition, information in such an organization must reflect the needs of managers. For example, information must be developed to allow the manager of accounting to control the accounting department, the marketing manager to control marketing, etc. In these situations the individual manager has substantial direct knowledge about the individual functions. Those needs are likely to be different from those of the manager who has little knowledge of the specific activity, as is the case in product form.

Markets: Information Requirement Issues

Instead of organizational structure, markets may also be used to co-ordinate multiple agents. In a pure decentralized market each buyer communicates directly with each seller. In a centralized market sellers work through brokers, who are in direct contact with the buyers. The information requirements of these markets can differ markedly because of the intermediary brokers. In a pure decentralized market, sellers need information about all buyers and buyers about sellers. How that information may be captured is also likely to vary from one market to another.

The addition of brokers to the market changes the information requirements. No longer do individual sellers and buyers need complete sets of information. Instead, brokers capture and maintain some of that information.

Differences in Information Requirements

A comparison of the different forms of multiple-agent structures leads to the conclusion that each has different information system requirements. Thus, an information system to meet those needs would vary from structure to structure.

Integration of Generic Problem Type and Multi-agent Structure

This paper has argued that *a priori* requirements analysis for information systems supporting intelligent systems can provide substantial information needs, as illustrated in the field study. Those needs can be assessed by considering the theories on which the intelligent system and the information system are based. In particular, a substantial base of information requirements can be specified by accounting for the generic problem type and the nature of multi-agent systems. The structure of the multi-agent systems can be further specified by including the structure of multi-agent systems (e.g. functional). Thus, together, these two bases provide a powerful *a priori* information requirements analysis tool, and are summarized in Figure 2.

Generic Problem Type: Human versus Computer System

Although the information system discussed in this paper is designed to support an intelligent diagnostic computer system, the discussion is not limited to support of computer systems. Instead, the intelligent diagnostic system may be a single agent or a set of multiple agents. Similarly, the support of generic problem type may be for either a human or a computer-based agent.

Summary and Extensions

This paper has described the development of an information system to support an intelligent system for diagnostic decision making. The resulting expert system and supporting information system formed an integrated system design architecture for the solution of a diagnostic decision problem in a multiple-agent environment.

It was found that information requirements for the supporting information system could be generated using the information system requirements suggested by theories of diagnostic problems and multi-agent environments. This general approach to the development of requirements for integrated, intelligent and conventional systems can be extended to other types of decisions (for example, planning, monitoring, etc.) and other information systems. In particular, because of the basic differences between decisions we would expect requirements for information systems designed to support different types of decisions to vary. In addition, we would expect that the requirements could be elicited by considering the specific generic type of decisions being made.

Single-agent/multi-agent structure

Generic Task\	Single Agent ...	Product ...	Functional ...	Market (D/C)*
Control				
Debugging				
Design				
Diagnosis				
Instruction				
Interpret-ation				
Monitoring				
Planning				
Prediction				
Repair				

Figure 2 Summary of requirements analysis structure for information systems designed to support intelligent systems. (D/C: decentralized/centralized.)

The single-agent/multi-agent nature of the decision problem also can be extended. In particular, the nature of the structure used to co-ordinate the multiple agents provides a basis for the development of information requirements to support the intelligent system. The structures examined included: product hierarchy, functional hierarchy, centralized markets and decentralized markets.

The merger of the concepts of decision type and multi-agent structure provides a powerful requirements analysis tool for information systems supporting intelligent systems (either human or computerized).

Extensions

There are a number of extensions of this research. First, knowledge and resulting information about agents may include additional information such as values and attitudes of those agents. Second, rather than a 'passive' information system to co-ordinate agent behavior, a group decision-support system could be developed. In such a system the agents could play a more active role of co-ordination by, for example, voting on different issues. Third, the system discussed in this paper dealt with a problem that was relatively deterministic. Further research could emphasize more probabilistic processes. Fourth, this research has ignored decision complexity, which could have a major impact on the requirements analysis process.

Acknowlegements

An earlier version of this paper was presented at the Philadelphia TIMS/ORSA Meeting, October 1990. The authors would like to thank the participants for their comments. Additional work on the paper by the first author was done while visiting the School of Computing Sciences at Bond University, Australia.

References

Bond, A. and Gasser, L. (1988) 'An analysis of problems and research in DAI', in *Readings in Distributed Artificial Intelligence*, San Mateo, CA: Morgan Kaufman.

Buchanan, B. and Shortliffe, E. (1985) *Rule-Based Expert Systems*, Reading, MA: Addison-Wesley.

Duda, R., Gaschnig, J. and Hart, P. (1979) 'Model design in the prospector consultant system for mineral exploration', in Mitchie, D. (ed.), *Expert Systems for the Micro Age*, Edinburgh: Edinburgh University Press.

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

Kaplan, A. (1964) *The Conduct of Inquiry*, San Francisco: Chandler Publishing.

Malone, T. (1987) 'Modeling coordination in organizations and markets', *Management Science*, **33**, No. 10, 1317–32.

Meservy, R. (1985) *Auditing Internal Controls: A Computational Model of the Review Process*, Unpublished PhD dissertation, University of Minnesota.

O'Leary, D. (1990) 'The security of expert systems', *IEEE Expert*.

O'Leary, D. and Watkins, P. (1990) 'Knowledge acquisition for a diagnostic-based task', in Schader, M. and Gant, W. (eds), NATO ASI Series, Vol. F61, *Knowledge, Data, and Computer-Assisted Decisions*, Berlin: Springer-Verlag, pp. 361–8.

Rockart, J. (1979) 'Chief executives define their own information systems needs', *Harvard Business Review*, March–April, 81–93.

Turban, E. and Watkins, P. (1986) 'Integrating expert systems and decision support systems', *MIS Quarterly*, June.

Watkins, P.R. and O'Leary, D. (1987) 'Knowledge acquisition for small scale expert systems from consulting experts: a field study approach', *Proceedings of the First European Workshop on Knowledge Acquisition for Knowledge-based Systems*, September, Reading University, UK.

Wetherbe, J. (1991) 'Executive information requirements: getting it right', *MIS Quarterly*, March, 51–65.

Williamson, O. (1970) *Corporate Control and Business Behavior*, Englewood Cliffs, NJ: Prentice-Hall.