Artificial Intelligence in Business, II: Development, Integration and Organizational Issues

Daniel E. O'Leary
School of Business
University of Southern California

John Kingston
AIAI
University of Edinburgh

An earlier version of this paper was presented at the Workshop on AI in Business at the IEEE Conference on Artificial Intelligence Applications, Orlando, Florida, March 1993.

Abstract

The purpose of this paper is to review the use of knowledge-based systems and artificial intelligence (AI) in business. Part I of this paper provided a broad survey of the use of AI in business, summarising the applications of AI in a number of business domains. In addition, it also provided a summary of the use of different forms of knowledge representation in business applications. Part I has a large set of references, including a number of survey papers, focusing on AI in business.

Part II of this paper consists of more detailed analysis of particular systems or issues affecting AI in business. It examines technical issues which are central to the construction of business AI systems, and it also examines the commercial contribution made by methods for the development of AI systems. In addition, part II looks at integration between AI and more traditional information systems. AI can be used to add value to many existing information systems, such as database management systems. Particular attention is given to the integration of AI with operations research, which is the one of the primary "competitors" of AI, providing an alternative set of support tools for decision making.

Business organizations are not concerned only with technology issues; there is also concern about the impact of AI on organizations. Further, the evaluation of AI often is based on an economic view of the world. Part II therefore investigates the organizational impact of AI, and the economics of AI, including issues such as value creation.

The format of Part II is as follows: Section 8 analyses techniques for improving the performance of AI systems, thus maximizing economic return.

Section 9 looks at different forms of uncertainty and ambiguity which must be dealt with by AI systems. It examines the contributions of fuzzy logic and numerical measures of certainty to handling these problems. Section 10 examines the usefulness of different approaches to knowledge acquisition in business situations, and investigates the benefits of methodological approaches to AI applications. It also looks at more recent AI programming techniques which eliminate the need for knowledge elicitation from an expert: neural networks, case based reasoning and genetic algorithms are discussed.

Sections 11 and 12 examine issues of integrating AI systems. Generally, the use of AI in business settings must ultimately be integrated with the broader base of corporate information systems. Chapter 11 looks at integration with information systems in general, and chapter 12 looks particularly at integration with operations research. Sections 13 and 14 review the organizational and economic impact of AI. Finally, section 15 provides a brief summary of part II.

8 Improving the performance of AI systems

The goal of all areas of business and industry in a competitive environment can be summed up in three words: *to work better*. Business processes are characterized by the need to make better **decisions**, rather than better products; many areas of business are highly competitive, with reputations or fortunes being won or lost on the basis of decision making skill. In many cases, today's businesses suffer from information overload, which means that the ability to make *faster* decisions allows more time for deeper analysis of the available information, thus contributing directly to an improvement in the quality of decisions.

The use of automation in the business world has always been targetted at the twin aims of better and faster decisions. Databases and spreadsheets have provided better access to information for decision makers; Artificial Intelligence promises to go further, by taking over some or all of the decision making process. The AI technology which has proved most successful at fulfilling this promise is *expert systems*. There are many hundreds of expert systems in use in the business community today; these systems have enabled the knowledge of key experts to be recorded, distributed and archived, with financial benefits sometimes running into millions of dollars [Feigenbaum *et al*, 1989].

This section discusses techniques for making rule-based expert systems produce good decisions in a shorter time. Other techniques for helping expert systems provide better decisions include better acquisition of knowledge (see section 10) or more accurate reasoning with the knowledge (section 9 examines a particular technical issue, that of reasoning with uncertain knowledge). Alternative AI programming techniques have also been proposed, and are beginning to find acceptance in the commercial world; one such technique is the use of *genetic algorithms*, which aim to find near-optimal solutions to highly constrained problems. Genetic algorithms are discussed in section 10.4.

8.1 Improving the performance of expert systems by ordering rules

In a rule-based expert system, the structure of the knowledge base and the inference engine's approach to processing that knowledge have a major impact on the information solicited and the order in which solutions are generated. The inferencing process controls both the amount of information solicited from the user, and which solutions are found first, by controlling the order in which rules are fired. Oftentimes decisions must be made in real-time and there are many feasible solutions; as a result, there is incentive to generate the best solutions as soon as possible. Unfortunately, the business representative may be able to present only one solution at a time to clients; and the presentation of one potential solution may preclude other solutions, on the basis of the time and effort required to produce further solutions, or to solicit additional information. The task of personal financial planning [Feigenbaum *et al*, 1989] is a good example of this. As a result, in many complex economic-based decisions, it is not surprising that customers are presented with sub-optimal solutions.

Improving the performance of expert systems is often dependent on careful organization of the knowledge base as much as acquisition of further knowledge. One approach to this organization is to establish partial orderings between the events or knowledge representation of the events (e.g., [Dean & Boddy, 1988]). In rule-based systems, this control can be achieved (depending on the inference engine) by the order of the rules. The "Bartender" problem described in [Winston, 1984] provides a clear example of how the order of the rules and conditions provided to the inference engine impacts decision making choices and potential returns to the firm and its agents.

One of the most important aspects of human decision making is the determination that one condition has a dominating effect, to the extent that other conditions are considered irrelevant (e.g., [Hogarth, 1985]). For example, life assurance companies which adjust premiums dependent on an applicant's state of health might consider a terminal disease to have a dominating effect over their decision; other diseases which the applicant might have would simply not be considered. In situations with dominating conditions, there may be a partial ordering of the rules and conditions such that those rules with the dominating condition should precede those rules that do not contain the dominating condition, in order to generate a good solution. Using this insight, [O'Leary & Watkins, 1992] finds that at least as good a set of solutions will be found if rules with a dominating condition are ordered before rules without that dominating condition. This approach forces the system to consider those rules with the dominating condition, before any other rules are considered. This and other results can be used to find orders of rules that provide better solutions than other approaches; and, for most businesses, better solutions contribute greatly to maximizing income.

9 Uncertainty and ambiguity in knowledge

Uncertainty of various forms has been a problem for AI systems since their earliest commercial applications [Johnson & Keravnou, 1985]. Many of the original techniques proposed for dealing with uncertainty have proved unsatisfactory. This is perhaps because these techniques attempted to use a single value to represent uncertainty, whereas there are (at least) five categories of uncertainty which affect the use of AI systems. Each of these five categories can lead to different interpretations being placed on the results of an AI system by different users; in most businesses, where there are many users of AI systems with different levels of skill, this is highly undesirable. The five categories are:

- fuzzy knowledge or data;
- ambiguous knowledge;
- lack of confidence in the knowledge or data;
- incomplete data;
- inconsistent knowledge or data.

These five categories are described and discussed below.

9.1 Fuzzy knowledge or data

Fuzzy knowledge is knowledge which is specified using terms such as *about*, *between*, or *approximately* (e.g. "about 10 per cent", "between 11:00 and 11:30"). It is also known as *imprecise* knowledge. Fuzzy knowledge is used in business applications worldwide for tasks such as retail pricing and predicting stock market returns (see [Hiemstra, 1993] and [Harmon, 1993a]); it is particularly common in Japan, where it is used in a wide variety of consumer products as well as for applications such as stock selection [Harmon, 1993a].

The best approaches for dealing with fuzzy knowledge are *fuzzy logic* and *fuzzy set theory* [Gaines, 1976] [Schmucker, 1984]. These approaches allow the programmer to represent the concepts of partial membership of a category, or partial matching against a pattern, by defining functions to calculate the goodness of fit between a datum and a range of values. Fuzzy techniques also work in reverse, allowing matching of precise values against a fuzzy set of requirements, or range of data, provided by the user of the system. An example of an application which uses fuzzy techniques can be found in [Chung & Inder, 1991], where techniques were developed to access geological databases containing fuzzy values, in order to support the task of petroleum exploration.

9.2 Ambiguity in knowledge

While most AI computer programming languages enforce the removal of much of the syntactic ambiguity in the acquired knowledge, semantic ambiguity occurs frequently in expert systems (e.g. [Szolovits, 1983]; [Hansen & Messier, 1986] ; [Steinbart, 1985]). For example, the following rule ([Buchanan & Shortliffe, 1984], p.512) from a financial planning expert system, requires the user to provide substantial financial information. Some of that information may differ from user to user, based on the interpretation provided by that user.

```
If   1) the client's income tax bracket is 50%, and
     2) the market has followed an upward trend recently, and
     3) the client manages his assets carefully
Then There is evidence that the area of the investment should be
     high technology.
```

In this example, the user is required to categorize evidence on many issues. Yet it is often unclear whether the market has followed an "upward trend" and whether or not that has occurred "recently." For example, what does the term "recently" refer to – the last week, the last month, the last year, etc.? Judgement is also required to decide if a client manages his assets "carefully" – but who is authorized to make such a judgement?

The extent to which semantic ambiguity is present in some business expert systems has been investigated by [O'Leary, 1990] [O'Leary, 1993a]. In tests of students and business consultants, given the same situations with evidence to be inputted to an expert system, different subjects gave opposite interpretations of the same data. It might be argued that experience or expertise can act to mitigate the impact of ambiguity in judgments. However, empirical studies ([O'Leary, 1990] [O'Leary, 1993a]), find that even expert judges are likely to come to different conclusions, given semantic ambiguity.

Semantic ambiguity remains a difficulty for AI systems; there is no one technique which can handle all cases of it. The difficulty appears to lie in the fact that experts actually use ambiguous representations. In some cases, the expert can be asked to provide more rigorous definitions of ambiguous terms; these may lead to a single precise representation, perhaps in a deontic logic (a logic which can be used to express legislative conditions; see [Jones & Sergot, 1992]) or, more commonly, to a fuzzy representation of knowledge which can be handled using fuzzy techniques. In other cases, the expert may never have needed to formulate a precise definition of his ambiguous statement; the knowledge engineer is then left to choose whether to persuade the expert to produce a definition (which may be an interesting learning exercise for the expert as well as the knowledge engineer), or to retain the ambiguous statement in the knowledge. There are a few cases where ambiguity is *deliberately* inserted (such as the frequent use of the term *reasonable* in laws & statutes) because it is desirable that the user should be given latitude in the interpretation of the regulation. This is a major stumbling block for AI systems.

In financial expert systems we see the use of terms such as "increasing," whose ambiguity seems to derive from particular situations. In some situations that term is likely to have clear definitions. For example, the time series 5, 10, 15 is clearly "increasing." However, with a time series like 5, 15, 10 it is unclear if the series is "increasing": both the 15 and the 10 are greater than the 5, but the 10 is less than the 15. Thus, there is concern with whether the second series would be categorized as "increasing." It is possible that statistical techniques such as trend analysis would be able to resolve some of this ambiguity. Trend analysis has previously been implemented in an AI system for technical analysis of share price movements [Merry & Prettejohn, 1988].

9.3 Lack of confidence in the knowledge or data

The earliest attempts to represent uncertainty in AI systems [Johnson & Keravnou, 1985] were concerned with representing lack of confidence in medical knowledge. The lack of confidence came from two sources: lack of confidence in the data, and lack of confidence in a a conclusion, given certain data. The chosen solution was to use *certainty factors*, which represented confidence using a number between 1.0 (total confidence that this datum is true) to -1.0 (total confidence that this datum is false). The value 0 represented a complete lack of information about the datum. Certainty factors were propagated through rules using a simple algorithm.

Certainty factors have the advantage of simplicity, but are theoretically flawed (the order in which rules fire can affect the final certainty factor if self-referential rules are used; see [Buchanan & Shortliffe, 1984]). However, other numerical representations of certainty (e.g. Bayes' theorem, which calculates probabilities of hypotheses based on evidence) have proved to be usable for representing and propagating data and knowledge which is not known with complete confidence.

9.4 Incomplete data

In some business situations there may be incomplete information available to the user. For example, in financial systems there can be different levels of data aggregation. If the expert system assumes a more detailed level than is available to the user, then the user would have to use estimation methods to generate more detailed information.

The best technique for handling incomplete data is to acquire the missing data. If this is not possible, then a number of techniques which handle partial matching might be appropriate, such as fuzzy techniques, forward chaining rules, or neural networks.

9.5 Inconsistent knowledge or data

Inconsistent knowledge should be rooted out of an AI system during the verification & validation of the system. Some computerized techniques are available for identifying inconsistencies in rule-based systems, such as Nguyen's graph method [Nguyen, 1987]. Inconsistent data is a more difficult problem to deal with. For example, if a financial expert system requires input based on inventory costing, this could result in different users providing different responses to the same question. This is because there are different measurement methods for inventory costing, such as the inventory cost being based on the cost of the first inventory acquired, or the inventory cost being based on the last inventory acquired, which can result in substantially different costs. There is no easy way to handle such inconsistent data, except the obvious method of requiring the user to specify the data measurement method used in sufficient detail for the AI system to be able to treat the value as unambiguous.

9.6 Impact of Uncertainty

AI systems are developed for a number of reasons in business contexts, e.g., ensuring consistency of judgment, delegating tasks downward (capturing expertise in the system and having others use the program), and making better decisions. If the resulting system does not produce the same results in two identical situations, then it can hardly be said to satisfy any of those requirements. Uncertainty and ambiguity can therefore have a strong negative effect on AI systems, unless steps are taken to represent or eliminate these factors.

In addition, AI systems need to be validated, i.e. tested against the original source of expertise. If ambiguous judgments interfere with that testing process, then there is some question as to the quality of the system. Ambiguity therefore makes validation of an AI system difficult to achieve.

10 Support for knowledge acquisition and knowledge analysis

This section reviews the contributions of new approaches to the development of AI systems in business. The vast majority of AI systems in business today are expert systems; these have traditionally been developed using a "rapid prototyping" methodology, in which the results of an interview with an expert are encoded into a prototype expert system. This prototype is then shown to the expert for validation, and is used as a trigger for further expert knowledge. However, there has been a growing desire amongst users of expert systems, particularly in the business community, for a more methodological approach to development of expert systems. Such approaches provide standardization of system development, and documentation of decisions and content, with perceived benefits in maintainability of systems and reusability of development techniques.

In addition, there has been a lot of interest in techniques which reduce the problem of the "knowledge acquisition bottleneck", either by providing support for the process of knowledge elicitation from experts, or by using programming techniques which eliminate the need to elicit knowledge from an expert. This section therefore discusses a variety of techniques in three subsections:

- Methodological approaches to expert system development;
- Techniques for supporting knowledge elicitation from experts;
- Programming techniques which do not require knowledge elicitation.

10.1 Methodological approaches to expert system development

While knowledge acquisition has presented a significant methodological problem to expert systems developers, the concept of "methodology" in the business world usually refers to approaches such as SSADM or Yourdon, which are focused on the analysis and design of a system. Equivalent methodological approaches have only made a significant impact in the knowledge engineering community in the last few years. These methods attempt to guide and to structure the development of expert systems and other knowledge-based systems, in order to speed up the development process, to provide justification and documentation for decisions taken, and (in a few cases) to produce output which is very close to being executable code.

In Western Europe, the most prominent methodological approach is the KADS methodology [Wielinga *et al*, 1992]. The KADS methodology views the development of knowledge based systems as a modelling process: models of various aspects of the knowledge are developed during the process of expert system development.

KADS provides extensive support for knowledge analysis from various viewpoints, and it is beginning to provide support for the system design phase as well. The main reason for the success of KADS is probably its library of generic inference models; if the *task type* (e.g. diagnosis, configuration) of a knowledge based system can be identified, then an inference model can be used to specify knowledge which is expected to be acquired, and to guide the structuring of that knowledge. KADS has already proved its worth in several major commercial knowledge-based projects, such as the identification of potentially fraudulent use of credit cards [Porter, 1992], and aircraft wing design [Bechtel AI Software, personal communication].

On the other side of the Atlantic, there is more focus on commercially-sponsored methods which start with acquired knowledge and transform the knowledge through a number of stages to produce executable code in a specified AI toolkit. The two most prominent methods are that promoted by Inference, which produces code suitable for implementation in ART Enterprise and that promoted by Trinzic (formerly Aion), which produces code suitable for implementation in ADS. For a survey of knowledge engineering methods and accompanying tools, see [Inder & Filby, 1991].

The business benefits provided by these methodological approaches include better documentation of the acquired knowledge, better documentation of decisions taken (in KADS, for example, the generation of models of various aspects of the knowledge requires modelling decisions to be made; these reflect decisions taken about knowledge structuring), and hence greater maintainability of the resulting system. The business benefits of improved documentation of knowledge are potentially vast; for example, KADS was used successfully to identify flaws in a business process by separately modelling the performance of experts, the performance of novices, and the approach recommended by the training manual [Bechtel AI Software, personal communication]. Good documentation also allows re-usability of previous approaches; KADS has proved that there is benefit in using generic inference models, and some of the current work on KADS is focusing on producing an architecture for reusable domain models.

Some methods attempt to speed up the process of expert system development; those methods which provide executable code are an obvious example. However, the managers of expert systems projects should be aware that the use of a methodology is likely to make small and medium-sized projects take *longer* than a rapid prototyping approach would require. The benefits are found in the reduced time taken at later phases of development, and at greatly reduced times for maintenance of the system. There is also less danger of the system growing in an uncontrolled mannner until it becomes unmanageable (see [Hart, 1988] for an example).

10.2 Techniques for supporting knowledge elicitation from experts

The task of knowledge acquisition has often been described as the chief bottleneck in the development of expert systems. Traditionally, knowledge acquisition has been conducted by interviews, requiring the knowledge engineer to transcribe and analyse the interview in order to obtain the requisite knowledge. Transcript analysis provides a useful record of the source of acquired knowledge, particularly if it is carried out on-line with hypertext links built between the transcript and the items of knowledge (cf. [Anjewierden *et al*, 1989]. However, even when using a computerized support tool, transcript analysis is time-consuming, prone to generate much irrelevant information, and provides no guarantees about the completeness of the knowledge acquired. Using structured interviewing techniques and protocol analysis (asking the expert to describe a particular case of problem-solving) solves some of these problems, but by no means all.

In order to overcome some of these problems, knowledge engineers have drawn on the field of psychology, and particularly on psychometric testing, in order to find other techniques which can be used to acquire knowledge. Some of the techniques which have proved particularly useful in expert systems projects are:

- **Card sort**. This is an apparently simple but surprisingly effective technique in which an expert categorizes cards which represent terms from the knowledge domain [Shadbolt & Burton, 1990];

- **Repertory grid**. This is a technique derived from psychotherapy in which an expert distinguishes objects in the domain; these distinctions are then analysed statistically to see if there is any implicit categorization of objects [Shaw & Gaines, 1987];

- **Laddered grid**. This technique uses key questions to persuade an expert to expand a taxonomic hierarchy to its fullest extent. By using different prompt questions, the laddered grid can also be used to elicit a simple procedural flow chart [Burton *et al*, 1988].

The benefits of these techniques are that they provide output in the form of categorizations, relationships and procedures, which resemble the formalisms required by expert system programming tools much more closely than a textual transcript does; they ensure consistent and fairly complete coverage in knowledge acquisition, by continual prompting (laddered grid) or by requiring all items to be categorized (card sort & repertory grid); and they are relatively simple to administer (although analysis of a repertory grid usually requires computer support). The repertory grid has been particularly well used, with over 150 applications to date having used this technique successfully [Boose, 1989].

10.3 Programming techniques which do not require knowledge elicitation

The emergence of AI programming techniques which do not require knowledge elicitation has been one of the most significant changes in the commercial AI world in the last 5 years. The techniques of **case based reasoning**, **rule induction**, and programming with **neural networks** can all be categorized as *machine learning* techniques, because their knowledge is based on automatic analysis of previous cases of problem solving. The techniques differ in the degree to which the previous cases are analysed and compiled before being used for problem solving; case based reasoning performs almost no compilation, rule induction compiles key factors into rules, and neural networks compile every factor which appears to be relevant into a statistically weighted network. Neural networks and case based reasoning have both proved popular in the business world, and are discussed below.

Another AI programming technique which avoids the need for knowledge elicitation is the use of **genetic algorithms**, which employ an approach based on the theory of evolution to search for an optimal solution to a problem. Genetic algorithms are also discussed below.

10.3.1 Neural networks in business

Although the techniques on which neural networks are based have been known for some decades, neural networks have only recently become established as a viable commercial technology. Neural networks identify and represent patterns of data, using a programming technique inspired by the physiology of the biological nervous system. This is accomplished by presenting a neural net with a "training set" of example cases, which it repeatedly analyses in order to extract patterns of data. The network can then be applied to further cases of the same type, in which it should be able to identify the same patterns and thus discriminate cases.

From the viewpoint of the business world, neural networks have two key attributes. The first is that they can perform pattern matching tasks on very large amounts of data efficiently. The second is that they "learn" patterns from a training set, and therefore do not require much programming ability in order to develop them. These two features have led neural networks to be applied to various business tasks, such as the following:

- Information Filtering. Neural networks can be used to screen large amounts of data and to classify the data into various categories, highlighting only those categories which are likely to be of interest to the user. The data can be as diverse as stocks, credit card transactions, news stories, or marketing databases (see [Harmon, 1992]). For example, Fidelity Investments are using neural networks to scan a universe of 2,000 stocks in order to finds stocks which are undervalued relative to their industries; the system suggests about

200 stocks for further consideration. The resulting fund has outperformed the market every year since its inception [Loufbourrow, 1992]. While expert systems are also well suited to classification tasks, neural networks have the advantage that they can be implemented without performing large amounts of knowledge acquisition. Instead, the criteria for the decision are automatically "learnt" from the training set; indeed, the Fidelity system is retrained every day.

- Forecasting. The task of forecasting the economic future is important to all businesses; for some, such as stock market trading, it is the very essence of the business. Many approaches to forecasting work by trying to identify economic conditions which are similar to those prevailing at some time in the past. This is clearly a complex pattern matching problem, and therefore one for which neural networks are well suited. Examples of neural networks which are used for forecasting include the prediction of IBM daily stock returns [White, 1988] and the SENN system, a commercially available tool for financial analysts which performs rate forecasting, portfolio management and risk analysis [**?**]

- Character recognition. Much of the work of most businesses involves the analysis and transfer of documents. Any computer system which can speed up this process will make a significant contribution to the profitability of the business. Neural networks can perform character recognition well, because they are able to identify the salient features of characters, and thus to perform correct discrimination even when the input is of very poor quality. For example, AT&T have developed a neural network which recognizes digits, in order to enable automatic processing of zip codes [LeCun & *et al*, 1989].

- Fault diagnosis. Neural networks, in common with many expert systems, are particularly useful for tasks involving troubleshooting or fault diagnosis. This is because fault diagnosis usually consists of identifying a fault by observing a particular pattern of symptoms, which corresponds well with the ability of neural networks to "learn" to recognise patterns of data.

It can be seen that neural networks have been used to tackle some problems which have previously been successfully solved using expert systems (such as information filtering and fault diagnosis) as well as other problems which have not previously been automated with any degree of commercial success.

Given the discussion earlier on ordering rules for maximum quality of decisions, it might well be asked if any similar technique can be applied to neural networks. However, the essence of neural networks is that they compile patterns into stimulus-response connections; there is little that a programmer can do to alter their behaviour. This feature is both a major strength and a major weakness of neural networks; while neural networks can be developed by someone with very little programming knowledge, it is difficult to be sure that the network is performing correctly, and very difficult to repair any flaws which are found. The training set is the key to the performance of the neural net; if the training set does not represent all the distinctions which are found in the domain, or also represents some coincidental distinctions, then the neural network may not function correctly – and the source of the problem may be very difficult to locate [Stader, 1992]. Neural networks are therefore most suitable for problems where very large training sets of data are readily available.

10.3.2 Case based reasoning

Case based reasoning avoids some the problems experienced by neural networks by reasoning with past cases of problem solving directly, instead of extracting patterns from the cases and compiling these into a network. In case based reasoning, a database of past cases of problem solving is built up. When a problem arises, the salient features of the new case are matched against the database of past cases, and the best match from the database is presented to the user. By basing reasoning on individual case histories, case based reasoning provides excellent opportunities for offering explanations to the user, as well as providing understandable reasoning.

There are currently few commercial applications of case based reasoning, most of which have been for manufacturing diagnostics, configuration of parts [Hennessey & Hinkle, 1991] and other manufacturing-related areas. A particularly interesting diagnostic application is Compaq's QUICKSOURCE system, which helps customers diagnose faults in Compaq printers without needing to call Compaq's customer support center [Nguyen *et al*, 1993] [Acorn & Walden, 1992]. However, ongoing research into the use of case-based reasoning for tutoring [Ashley & Aleven, 1991] and for legal reasoning [Skalak & Rissland, 1992] [Sanders, 1991] suggests that this technology may well provide significant business-related applications in the future.

10.4 Genetic algorithms

Genetic algorithms attempt to "evolve" a solution to a problem by repeatedly altering existing solutions and evaluating the results. They are an even more recent phenomenon in the business world than neural nets, which means that their degree of usefulness is not yet clear. However, they have proved their worth in the production of near-optimal solutions to highly constrained peoblems which have so many possible solutions that searching for the optimal solution is not feasible. A good example of such a problem is the task of timetabling; the scheduling of timetables is known to be a difficult problem, both in business [Cadas, 1989] and elsewhere. An example of the use of a genetic algorithm for timetable scheduling can be found in [Fang *et al*, 1993], in which the genetic algorithm produced a much better schedule than was produced by hand, because the human scheduler was unable to cope with the number of constraints on the task.

The main interest in the use of genetic algorithms in business currently lies in hybrid applications, where genetic algorithms are used in conjunction with neural nets, expert systems, or other AI techniques. This approach was used for the FX Trader system developed by Citibank, where genetic algorithms are used to identify the most accurate technical indicators for forecasting foreign exchange rate changes, and the resulting indicators are used as input to a neural network [Colin, 1992] [Loufbourrow, 1992].

11 Integrating AI and Information Systems

AI comes to the business environment facing the need to integrate with information systems (IS) departments. Generally, IS controls computing throughout most business organizations. It is only through IS that there is access to databases or even computing capabilities. As a result, there are a number of incentives to integrate AI with information systems. A general survey of integration issues is presented in [Watkins & O'Leary, 1993].

Although AI systems and information systems must be closely coupled in virtually all business situations, it is only in the last 2 or 3 years that reports of commercial AI systems have included a discussion of the integration of the AI system and the related information system. However, the reports which have emerged suggest that integration of AI systems is becoming widespread in many areas of IS. For example, the proceedings of the 1992 conference on Innovative Applications of Artificial Intelligence [Scott & Klahr, 1992] describe systems which have integrated AI with call logging systems, charting systems, relational databases, client-server architectures, computer-supported cooperative work technology, a C grammar parser, user interface building software, and COBOL database access.

[O'Leary & Watkins, 1992] suggest that there are two extreme types of integration of intelligent systems and traditional information systems. First, an intelligent system could be embedded in a conventional information system to support some set of activities. Second, a traditional information system might be developed to support an AI system.

11.1 Integration of AI with IS: Case Study

[O'Leary & Watkins, 1992] present a case study on the development of an expert system and the information system to support that system.

[O'Leary & Watkins, 1992] developed a diagnosis system for a company that provided income tax services. Virtually all the clients of the firm for which the system was to be built had micro computers with modems, in order to communicate with the firm's computer facilities.

The company therefore uses diagnosticians to provide a "help desk" service which assists clients in the process of communication and use of the software. All questions from clients are answered over the phone. Either questions are answered at the time of the call or the diagnostician calls the client back. Sometimes this is an iterative process, with the diagnostician offering a solution, getting feedback from the client and then generating another solution. Generally, calls are handled by available agents, although there is an effort to use the agent who has dealt with the client in the past.

In order to provide an appropriate response to the client, the diagnostician has both knowledge and information needs that must be coordinated. In addition, since multiple diagnosticians are often involved in responding to a single client, the multiple agents must be coordinated.

A rule-based expert system was built to include the expertise of the more expert agents in diagnosing client problems with modems. In addition, an information system was integrated with the system to support the use of the expert system. The system architecture is summarized in figure 1.

```
-------------------------
Figure 1 about here
-------------------------
```

The integrated system has a knowledge base of symptoms, problems and solutions and an inference engine, as in a classic rule-based system. In addition, the system is coupled to the firm's database of customer information, which includes the client contact and the type of hardware and software at the specific client; and the system captures the client's history, based on changes that occur. In this application the access to client information and the addition of new information as it occurs are as important to problem diagnosis as the underlying knowledge used to make the diagnosis.

11.2 Integrating AI and Databases

One of the key applications of AI in the business world has been that of *data mining* – extracting information from the huge databases which some businesses have accumulated since the advent of database technology in the 1960s. However, this key application has been stymied for many years because of the incompatibility between the languages and the hardware of AI systems and databases. Only in the last 5-10 years has integration between AI systems and corporate database become a feasible option.

Initially, integration focused on either integrating AI into numeric accounting-like databases or into text-based databases; [O'Leary, 1991] provides a survey of the literature. For example, [Arthur Anderson, 1985b] [Arthur Anderson, 1985a] and [Mui & McCarthy, 1987] discuss two AI systems developed by Arthur Anderson to interface with EDGAR, a large database belonging to the US Security and Exchange Commission. EDGAR was designed to hold accounting information from all publicly held companies in the USA. The first AI system, ELOISE [Arthur Anderson, 1985b] was designed to search through an ASCII database in order to find information that related to a specific acquisition issue "anti-takeover provisions." The second system, FSA, was designed to search through various disclosures (also presented in ASCII) in order to gather information to calculate various financial ratios. These systems employed the work of DeJong (1979) to structure the understanding of text data.

More recently, the advance of technology and the increasing competitiveness of the business world has led to more widespread use of AI techniques to access dtatabases. The advent of neural networks has opened up a new avenue into large databases; for example, Advanced Software Applications Corporation of Pittsburgh, PA, now offer a commercially available neural-net based package for weeding out likely prospects for direct mailings from large marketing databases [Harmon, 1993b]. Also, significant work has been done on extracting information from gigabyte-sized databases; a good example appears in [Anand & Kahn, 1992], in which marketing information is extracted from a database of point-of-sale information much faster than any marketing analyst was able to achieve.

Decision support systems which support decision making by providing the user access to database systems and to tools to analyze those databases have also been enhanced using expert systems capabilities. [Turbin & Watkins, 1986] provided perhaps the first assessment of integration issues between these types of systems.

12 Integrating AI and Operations Research

One of the primary "competitors" of AI in business is operations research (OR, also know as management science or as operational research). Generally, operations research is viewed as the application of a variety of mathematical and systematic approaches to the solution of a set of business problems. Among those approaches are mathematical programming (such as constrained linear optimization theory – linear programming), network analysis and queueing theory. Operations research tools are used to generate optimal solutions to the constructed problems.

OR is often considered to be a competitor of AI, which suggests that AI and OR are interested in solving many of the same problems. OR has different strengths and weaknesses from AI: the formulation and ultimate interpretation of OR algorithms is often too costly or time consuming to impact decision making; the use of many OR approaches requires substantial expertise in OR and in the specific software, particularly if the OR models are hard-coded to improve performance; and

OR methods provide a single value or set of values as a solution, based on input values which are assumed to be correct, which is too simplistic an approach for a complex multi-process system. As a result, as noted by [Fabozzi & Valente, 1976], oftentimes OR is not used. Thus, researchers have investigated the use of AI to make OR accessible in a timely and cost-effective manner, to generate more powerful problem-solving approaches. A recent survey of some of those efforts is provided in [McBride & O'Leary, 1993].

12.1 Using AI to formulate & interpret OR

One system that integrates AI and OR is CASHMANAGER [McBride *et al*, 1989]). CASHMANAGER is a system designed to support a corporate user in the process of making decisions about which financial instruments should be choosen as cash investments. The user faces a very complex (and generally a real time) decision of whether to invest in a given financial instrument, given an existing portfolio of other financial instruments investments. The choice among instruments can be made using an OR approach, linear programming. However, it is difficult and time consuming to formulate and interpret the solution of the output using that OR approach. As a result, linear programming is not often used to solve cash management problems, even though it does give an optimal solution to the problem.

CASHMANAGER uses object representations of financial instruments. Those instruments, such as a short term loan, are captured as representations of networks, with cash inflows and outflows occurring over time. Thus, when the user chooses a set of instruments CASHMANAGER assembles an aggregated network of those instruments and uses that network to formulate a linear programming representation of the problem. The system then proceeds to solve the resulting representation and provides the user with an interpretation of the solution. CASHMANAGER embeds the intelligence of an operations research expert in the context of the system.

CASHMANAGER also includes substantial knowledge of instruments, in the form of an interfaced rule-based expert system. In particular, the system contains "common sense" about the instruments so that it makes suggestions about which instruments should or should not be coupled together. For example, the system considers the state of the economy and the riskiness of individual instruments before suggesting that some instrument should not be considered in the linear program. This is important since the OR representation does not consider those factors, yet they are important to decision makers.

12.2 Using AI to apply OR to complex problems

The Quality and Reliability Expert System (QRES) [Moon *et al*, 1993] developed by IBM is a hybrid AI/OR system for quality management applications. QRES is used to identify problem areas in manufacturing production lines; to identify where to make improvements to maximize quality levels for the least cost; and to predict production quality levels for new components and assemblies during design and development. IBM had already tried a pure OR approach to tackle quality and reliability problems, but it was found that a complex problem domain such as quality management could not be solved with a single solution. Instead, OR tools were used in narrowly focused areas which were then combined and extrapolated by experts.

The QRES system used AI technology to link the various OR modules, by handling the complex task of defect propagation from one OR module to the next. The production environment is already described in computerized models, so forward chaining rules are used to deduce the expected quality levels at each step. For reliability prediction, a backward chaining approach is incorporated in order to determine which step of production is responsible for known defects. Products, sectors and defects are represented using an object-oriented representation.

QRES has a wide range of uses across many organizations: within IBM, it is being used at 4 different sites, in "user organizations" ranging from Quality engineering and Test engineering through to Reliability and Service.

13 The Impact of AI on Organizations

An area of research getting increasing attention is the impact of AI on organizations. [Duchessi *et al*, 1993] provide a survey paper summarizing many of the issues in this evolving area of research. One of the more comprehensive studies of a single organization is developed in [Trewin, 1991]. The objective of that research was to identify the changes that occurred in a professional accounting firm with the introduction of an expert system. The research was limited to a single firm (Coopers & Lybrand) and a single expert system (ExperTax). The system has been discussed in the literature by [Shpilberg & Graham, 1986], [Shpilberg *et al*, 1986] and [Sviokla, 1990].

13.1 Case Study

Research on the impact of AI in organizations can take any of a number of directions. One approach is to survey or interview individuals from multiple different organizations. This approach can be used to generate a number of different data points. An alternative approach is to focus on a single organization, generating depth at the cost of a number of data points.

In addition, the location of the conduct of the research might vary. In the case of a survey approach, the research would not need to visit the actual locations in which the system is used. However, in a field study, the researcher would need to visit locations in which the system is used.

The approach taken by [Trewin, 1991] was to focus on multiple locations of the same organization. A field study approach was used.

[Trewin, 1991] found that the politics of the individual offices did not change after the introduction of the system. However, the introduction of the system did result in some changes in the technology, the structure of the firm and in the culture of the individual offices of the firm, throughout the country. In addition, she also found that it was not the system alone that forced the impact; guidelines for the use of the system also had a major impact.

The use of ExperTax changed the method of data collection from a manual solitary task to a computerized task. In addition, the use of the system helped the decision makers gain additional insight into business processes and goals of clients.

The structure of the firm refers to the formal organization and the lines of authority and responsibility. ExperTax changed the responsibilities of those using the system. In particular, the collection of data was generally raised upward to the managerial level. In addition, the guidelines for the use of the system indicated that two departments (audit and tax) within the offices would be responsible for collecting the data. Prior to the introduction of the system, the tax department had only been responsible for reviewing the collected data.

The organization's culture (informal rules and relationships among individuals and groups) also changed. The system and guidelines for its use increased the formal and also informal interaction between the two departments, tax and audit. In addition, the use of the system made members of the firm proud that their firm had taken such a leadership role in the industry.

The system and its guidelines for use also affected job content, productivity, and quality. Job content at the data collection stage changed because of the computerization. In addition, job content changed because of the interaction of the two departments. Data collection productivity increased because of the computerization of the process. The quality of the responses in the data collection effort improved because of the level of the personnel involved and because both departments were involved in the process.

13.2 Using AI for Business Process Re-engineering

The organizational changes highlighted in the case study above are typical of the beneficial effects of reorganising a problematic business process, with an AI system as a key component in the reorganized process. Many organizations today desire to be able to monitor and improve on their performance against their strategic objectives. As a result, many consultancies are offering services in *business process re-engineering* – reorganising a business' processes in a manner which is both coherent and consistent with strategic objectives.

It has frequently been demonstrated that expert systems support organizations in capturing and modelling certain key processes through providing them with enriched modelling and representation techniques. It therefore seems likely that AI techniques could be extended to capture a large number of business processes, model the processes in a standard format, and make coherent plans for re-organization. Knowledge acquisition techniques could be used for capturing processes; AI model-based representations could be used to model the processes; and AI planning techniques could be used to plan the re-organization [Anderson, 1993].

A project is currently under way to provide a computer toolset to support all the above functionality, as well as simulation and evaluation of revised workflow management systems [Fraser, 1993]. Given the potential of AI techniques for supporting this area, it seems likely that, as automated tools become available to support business process re-engineering, AI features will be closely integrated with many of these tools.

14 Economics of AI

One of the primary strategies of business in today's world is to choose approaches that create value. In that sense the use of AI can be viewed as a process that should occur to the extent that it creates value. As a result, since business decisions are made considering value creation, it is critical that knowledge engineers understand the bases of management's decisions. ([O'Leary, 1993b])

Value is probably is most generally viewed as a the difference between the cost and benefits of a system. Although in many cases the costs can be readily identified, the benefits occur over a longer period of time and are more difficult to identify. Value creation in the firm is an issue that has received attention by researchers in the economics of the firm, as discussed in the economics of internal organization ([Williamson, 1975]), finance (e.g., [Fruhan, 1979]), industrial organization (e.g., [Bain, 1968]) and strategy (e.g., [Porter, 1980]). These contributions are summarized here as the economics of strategy.

14.1 Creation of Barriers to Entry

In the economics of value creation (e.g., [Fruhan, 1979]) one of the approaches toward developing value is to foster the creation of barriers to entry of other firms. As noted by [Fruhan, 1979] "Entry barriers make it possible for a firm to increase operating revenues above (or reduce operating cost below) levels that would otherwise exist in a fully competitive situation."

[Bain, 1968], p.255 lists some sources that function as barriers to entry. These barriers include, "Product differentiation advantages established over potential entrant firms" and "Absolute cost advantage of established over potential entrant firms." Similarly, [Porter, 1980] elicits what are referred to as three generic strategies: overall cost leadership (requires efficient facilities, vigorous pursuit of cost reductions and cost minimization), differentiation (something that is perceived in the industry as being unique) and focus (concentrating on a particular buyer, product line or geographic market). The first two are similar to those of [Bain, 1968]. Other such barriers might include quality or reliability.

AI can assist the firm in developing such barriers to entry. Cost leadership might be attained by automating jobs done by human workers using systems. Discussions with one executive identified an application that led to the development of a computer system and the elimination of a "room full of clerks" ([O'Leary & Watkins, 1992]). Now instead of those clerks, there is a system manager who remains to maintain the system. Thus, systems might also reduce costs to the point where a barrier to entry could be developed.

Cost leadership is not limited simply to reducing wages. Commercial loan decision systems (e.g., [Duchessi *et al*, 1988]) can assist in the automation of certain loan officer activity. As part of the analysis of loans, such systems typically are designed to minimize costs incurred, such as loans not repaid, and maximize interest received.

Computer systems also can function as a basis of product differentiation. For example, Peat Marwick's expert system "Loan Probe" ([Willingham & Ribar, 1988] and [Ribar, 1988]) was designed to assist in the analysis of the evaluation of the quality of loans of a financial institution. Peat Marwick already holds a large portion of the market for financial institutions. This system gives them some additional product differentiation from other audit firms since the use of the auditing system provides a formal structure to the audit not offered by other firms.

Product differentiation also can be attained with the use of systems designed to ensure security of a service. TRW's system "DISCOVERY" ([Tenor, 1988]) is the only system designed to monitor and secure a commercial credit history file. As a result, services rendered by the system (determining unusual client agent accesses – say at 3:00 AM on a remote printer) provide their clients with a unique service.

Further, systems can assist firms in focus. Some computer systems, such as DSS, are aimed at specific problems. These systems are narrowly defined in terms of purpose and function, in part, due to the technology and in part, due to the understanding brought to the capture of a problem not previously put in a computer environment. The system Loan Probe can be viewed as assisting in focusing efforts of Peat Marwick. Peat Marwick identified the banking industry as one of the industries in which they are specialized. Loan Probe was developed as system for the banking industry. The development of Loan Probe formally identified those within the firm that were viewed as experts,it committed additional resources to that industry, and it provided a very visible commitment to that industry.

Quality, reliability and speed also can create barriers to entry. The "Authorizor's Assistant" developed by American Express ([Davis, 1987]) provides the ability of that firm to respond to card member purchases in a timely manner, while providing a high quality of service.

14.2 Reducing Risk of Doing Business

Another approach suggested by the economics of value creation is the reduction of risk. As noted by [Fruhan, 1979], "A firm can sometimes ... reduce its business risk below that experienced by less imaginative competitors ..."

Computer systems allow a reduction of risk for a number of different reasons. First, computer systems can allow the firm to increase consistency of problem solving approaches ([Willingham & Ribar, 1988]). Such consistency can lead to a decrease in the variance of behaviors and a corresponding increase in quality. Consistency is particularly critical in financial systems, such as American Express's "Authorizor's Assistant," ([Newquist, 1987]) where lower level personnel are using the system to perform higher level activities.

Second, by documenting the decision process, these systems provide a record of the process, thus reducing the risk that there will be no such record of why decisions were made. In addition, the existence of documentation provides a basis on which to evaluate the actual risk. As noted by [Willingham & Ribar, 1988] in the discussion of an audit-based system, "Through the proper design of ... systems, the required documentation for a given audit judgment can be automatically provided as part of the output of the judgment exercise ...." Similar statements can be made for credit granting systems, security systems, etc.

14.3 First Mover Effects

First mover effects as a phenomenon have been described as follows by Williamson (see [Williamson, 1975]):

> Winners of initial contracts acquire, in a learning by doing fashion, nontrivial information advantages over nonwinners. Consequently, even though large-numbers competition may have been feasible at the time the initial award was made, parity no longer holds at the contract renewal interval. The information acquired through experience is impacted in the sense that (1) original winners may refuse to disclose it (which is a manifestation of opportunism) or (2) they may be unable, despite best efforts to disclose it (because of bounded rationality of the language impeded variety). Small numbers of bargaining situations thus evolve in this way."

Recently, one of the authors (DO'L) was involved in discussions with a large international financial organization that was pursuing the development of a system because it did not want its direct competitors to gain any additional advantage. The organization felt that if they were to wait until their competitors developed and deployed such technology that their competitors might have an insurmountable, first mover advantage.

15 Summary

Part I of this paper provided an overview of some of the unique aspects of AI in business settings. A broad base of applications was reviewed in a variety of domains, including auditing, credit and loan systems, internal auditing, management accounting, financial planning, tax and other areas. In addition, the paper provided an extensive set of references including survey papers in a number of those domains. The use of a number of forms of knowledge representation methodologies in business settings was also investigated, including rules, cases, Bayes' nets and multiple agent systems.

The implementation of AI in business settings has led to an number of innovations in both knowledge acquisition and explanation. Part I summarized some of those contributions. In addition, the first part of this paper discussed the relationship between normative and descriptive models in AI research in business. Finally, part I provided a survey of some additional research resources on the use of AI in busines.

Part II of this paper investigated a number of issues in greater detail. The first three sections investigated issues of concern in the development of expert systems: designing expert systems in order to provide better decisions faster; analysing and representing ambiguity in knowledge; and improving the process of developing AI systems by using psychology-based knowledge acquisition techniques, by adhering to a development methodology, or by avoiding knowledge elicitation altogether by using a machine learning approach.

Integration of AI with information systems and operational research is important. AI has been integrated with many information systems, particularly databases; it has also been integrated with OR. There is an increasing trend for AI to enter the business world via integration with existing systems, rather than as a stand-alone package.

Business organizations are not concerned only with the technology side of AI. Instead, there is interest in the organizational impact and economic impact of AI. AI can make a beneficial impact on an organization's processes, and could potentially be used to plan the re-engineering of those processes. It can also provide an organization with cost leadership, product differentiation, and focus.

References

[Acorn & Walden, 1992] Acorn, T.L. and Walden, S.H. (1992). SMART: Support Management Automated Reasoning technology for Compaq Customer Service. In A.C. Scott and P. Klahr, (ed.), *Innovative Applications of Artificial Intelligence 4*, pages 3–18. AAAI Press.

[Anand & Kahn, 1992] Anand, T. and Kahn, G. (1992). Making Sense of Gigabytes: A System for Knowledge-Based Market Analysis. In A.C. Scott and P. Klahr, (ed.), *Innovative Applications of Artificial Intelligence 4*, pages 57–69. AAAI Press.

[Anderson, 1993] Anderson, G., (ed.). (1993). *AIRING*, volume 16. AIAI, University of Edinburgh, Scotland.

[Anjewierden *et al*, 1989] Anjewierden, A., Wielemaker, J. and Toussaint, C. (1989). Shelley – Computer Aided Knowledge Engineering. In et al, Wielinga, (ed.), *Current Trends in Knowledge Acquisition*, pages 41–59. IOS Press.

[Arthur Anderson, 1985a] Arthur Anderson, (1985a). ELOISE Project Report. unpublished report.

[Arthur Anderson, 1985b] Arthur Anderson, (1985b). Financial Statement Analyzer. unpublished report.

[Ashley & Aleven, 1991] Ashley, K.D. and Aleven, V. (May 1991). A Computational Approach to Explaining Case-based Concepts of Relevance in a Tutorial Context. In *"Proceedings of the Case-based Reasoning Workshop"*, Washington DC. DARPA.

[Bain, 1968] Bain, J. (1968). *Industrial Organization*. John Wiley, New York.

[Boose, 1989] Boose, J.H. (1989). A survey of knowledge acquisition techniques and tools. *Knowledge Acquisition*, 1(1).

[Buchanan & Shortliffe, 1984] Buchanan, B. and Shortliffe, E.H. (1984). *Rule-based Expert Systems*. Addison Wesley.

[Burton *et al*, 1988] Burton, A.M., Shadbolt, N.R., Rugg, G. and Hedgecock, A.P. (1988). Knowledge Elicitation Techniques in Classification Domains. In *Proceedings of ECAI-88: The 8th European Conference on Artificial Intelligence*.

[Cadas, 1989] Cadas, C.N. (1989). Meeples – an expert system for scheduling meetings. In Shadbolt, N., (ed.), *Research and Development in Expert Systems VI*, pages 207–219. Cambridge University Press.

[Chung & Inder, 1991] Chung, P.W.H. and Inder, R. (August 1991). Handling Uncertainty in Accessing Petroleum Exploration Data. Technical Report AIAI-TR-95, AIAI, University of Edinburgh.

[Colin, 1992] Colin, A.M. (Nov 1-5 1992). Neural networks and genetic algorithms for exchange rate forecasting. In *Proceedings of the International Joint Conference on Neural Networks*, Beijing.

[Davis, 1987] Davis, D. (April 1987). Artificial intelligence goes to work. *High Technology*.

[Dean & Boddy, 1988] Dean, T. and Boddy, M. (October 1988). Reasoning About Partially Ordered Events. *Artificial Intelligence*, 36(3).

[Duchessi *et al*, 1988] Duchessi, P., Shawky, H. and Seagle, J. (1988). A Knowledge-engineered System for Commercial Loan Decisions. *Financial Management*, 17(3):57–65.

[Duchessi *et al*, 1993] Duchessi, P., O'Keefe, R., and O'Leary, D. (1993). A Research Perspective on Artificial Intelligence, Management and Organizations. *International Journal of Intelligent Systems in Accounting, Finance and Management*, 2(3).

[Fabozzi & Valente, 1976] Fabozzi, F. and Valente, J. (1976). Mathematical Programming in American Companies: A Sample Survey. *Interfaces*, 7:93–98.

[Fang *et al*, 1993] Fang, H., Corne, D. and Mellish, C. (June 1-4 1993). Solving the Modular Exam Scheduling Problem using Genetic Algorithms. In *Proceedings of the 6th International Conference on Industrial and Engineering Applications of Artificial Intelligence and Expert Systems*, Edinburgh. Also available from Department of Artificial Intelligence, University of Edinburgh.

[Feigenbaum *et al*, 1989] Feigenbaum, E., McCorduck, P. and Nii, P. (1989). *The Rise of the Expert Company*. McGraw-Hill.

[Fraser, 1993] Fraser, J. (Aug 1993). The Enterprise Project. *AIRING*, (16):6–8.

[Fruhan, 1979] Fruhan, W. (1979). *Financial Strategy.* R. D. Irwin, Homewood, Illinois.

[Gaines, 1976] Gaines, B.R. (1976). Foundations of fuzzy reasoning. *International Journal of Man-Machine Studies*, (8):623–668.

[Hansen & Messier, 1986] Hansen, J. and Messier, W. (Fall 1986). A Preliminary Investigation of EDP-XPERT. *Auditing: A Journal of Theory and Practice*, 6(1):374.

[Harmon, 1992] Harmon, P., (ed.). (Jul 1992). *Intelligent Software Strategies.* Number VIII, 7.

[Harmon, 1993a] Harmon, P., (ed.). (Apr 1993). *Intelligent Software Strategies.* Number IX, 4.

[Harmon, 1993b] Harmon, P., (ed.). (Sep 1993). *Intelligent Software Strategies.* Number IX, 9.

[Hart, 1988] Hart, A. (1988). *Expert Systems for Managers.* Kogan Page.

[Hennessey & Hinkle, 1991] Hennessey, D. and Hinkle, D. (May 1991). Initial Results from Clavier: A Case-based Autoclave Loading Assistant. In *Proceedings of the Case-based Reasoning Workshop*, Washington DC. DARPA.

[Hiemstra, 1993] Hiemstra, Y. (Aug 24-28 Aug 1993). A fuzzy logic model for the prediction of quarterly stock market excess returns. In *Submitted to the Third International Conference on Artificial Intelligence on Economics and Management*, Portland, Oregon.

[Hogarth, 1985] Hogarth, R. (1985). In *Judgement and Choice*, Chichester. Wiley.

[Inder & Filby, 1991] Inder, R. and Filby, I. (Dec 1991). A survey of knowledge engineering methods and supporting tools. In *Proceedings of the BCS SGES Workshop on Methodologies for Intelligent Systems*, RHS Conference Centre, London. Also available from AIAI as AIAI-TR-99.

[Johnson & Keravnou, 1985] Johnson, L. and Keravnou, E. (1985). *Expert Systems Technology.* Abacus Press.

[Jones & Sergot, 1992] Jones, A.J.I. and Sergot, M. (1992). Deontic Logic in the Representation of Law: Towards a Methodology. *Artificial Intelligence and Law*, (1(1)):45–64.

[LeCun & *et al*, 1989] LeCun, Y. *et al.* (1989). Backpropagation Applied to Hand Written Zip Code Recognition. *Neural Computation*, (1(4)):541–551.

[Loufbourrow, 1992] Loufbourrow, T. (1992). Advanced Trading Technologies: AI on Wall Street & World Wide. *Intelligent Software Strategies.*

[McBride & O'Leary, 1993] McBride, R. and O'Leary, D. (Forthcoming 1993). Integrating Artificial Intelligence and Mathematical Programming. *European Journal of Operational Research.*

[McBride *et al*, 1989] McBride, R., O'Leary, D. and Widmeyer, G. (1989). A Knowledge-based System for Cash Management Using Management Science Expertise. *Annals of Operations Research.*

[Merry & Prettejohn, 1988] Merry, M. and Prettejohn, M. (1988). EDWIN: A Knowledge-Based system for Chart Analysis. In *Proceedings of the 8th International conference on Expert Systems and their Applications*, number 2, pages 159–169, Avignon.

[Moon *et al*, 1993] Moon, C.S., Moore, T.M. and Nassar, S.M. (1993). Managing Product Quality by Integrating Operations Research and Artificial Intelligence Technologies. In P. Klahr and E. Byrnes, (ed.), *Proceedings of the Fifth Innovative Applications of Artificial Intelligence Conference*, pages 126–141. AAAI Press.

[Mui & McCarthy, 1987] Mui, C. and McCarthy, W. (1987). Applying AI Techniques to the Familiarization Phase of Decision Making. *IEEE Expert*, 2(3):33–41.

[Newquist, 1987] Newquist, H. (April 1987). American Express and AI: Don't Leave Home Without Them. *AI Expert*, pages 63–65.

[Nguyen, 1987] Nguyen, T.A. (1987). Verifying consistency of production systems. In *Proceedings of the Third Conference on AI Applications*, Washington DC. IEEE.

[Nguyen *et al*, 1993] Nguyen, T., Czerwinski, M. and Lee, D. (1993). COMPAQ QuickSource: Providing the Consumer with the Power of Artificial Intelligence. In P. Klahr and E. Byrnes, (ed.), *Innovative Applications of Artificial Intelligence 5*, pages 142–151. AAAI Press.

[O'Leary & Watkins, 1992] O'Leary, D. and Watkins, P. (May 1992). Integration of Intelligent Systems and Conventional Systems: Requirements for Coordinating Multiple Agents for Diagnostic Decisions. *International Journal of Intelligent Systems in Accounting, Finance and Management*, 1(2):135–150.

[O'Leary, 1990] O'Leary, D., (1990). The use of imprecise language in knowledge-based systems for auditing judgment. paper presented at the American Accounting Association National Meeting.

[O'Leary, 1991] O'Leary, D. (1991). Artificial intelligence in accounting databases. *Expert Systems with Applications*.

[O'Leary, 1993a] O'Leary, D., (1993a). The impact of semantic ambiguity on the dependability of expert systems. Unpublished working paper under review.

[O'Leary, 1993b] O'Leary, D. (1993b). Value creation from expert system: An economic approach with applications in accounting, auditing and finance. In *Artificial Intelligence in Economics and Business Administration*. forthcoming. Presented at the University of Pennsylvannia, Feb 1990.

[Porter, 1980] Porter, M. (1980). *Competitive Strategy*. Free Press, New York.

[Porter, 1992] Porter, D. (1992). Reusable analysis and design components for knowledge based system development. In Th. Wetter *et al*, (ed.), *Current Developments in Knowledge Acquisition: EKAW-92*. Springer-Verlag.

[Ribar, 1988] Ribar, G. (June 1988). Development of an expert system. *Expert Systems Review*, pages 3–8.

[Sanders, 1991] Sanders, K.E. (May 1991). Within the letter of the law: Reasoning among multiple cases. In *Proceedings of the Case-based Reasoning Workshop*, Washington DC. DARPA.

[Schmucker, 1984] Schmucker, K.J. (1984). *Fuzzy Sets, Natural Language Computations and Risk Analysis*. Computer Science Press.

[Scott & Klahr, 1992] Scott, A.C. and Klahr, P., (eds.). (1992). *Innovative Applications of Artificial Intelligence 4*. AAAI Press.

[Shadbolt & Burton, 1990] Shadbolt, N. and Burton, A.M. (1990). Knowledge elicitation. In J. Wilson and N. Corlett, (ed.), *Evaluation of Human Work: A Practical Ergonomics Methodology*, pages 321–346. Taylor and Francis.

[Shaw & Gaines, 1987] Shaw, M.L.G. and Gaines, B.R. (1987). An Interactive Knowledge Elicitation Technique using Personal Construct Technology. In Kidd, A., (ed.), *Knowledge Acquisition for Expert Systems: A Practical Handbook*. Plenum Press.

[Shpilberg & Graham, 1986] Shpilberg, D. and Graham, L. (Fall 1986). ExperTax: Developing ExperTax: An Expert System for Corporate Tax Accrual and Planning. *Auditing: A Journal of Theory and Practice*.

[Shpilberg *et al*, 1986] Shpilberg, D., L., Graham and Schatz, H. (July 1986). ExperTax: An Expert System for Corporate Tax Planning. *Expert Systems*.

[Skalak & Rissland, 1992] Skalak, D.B. and Rissland, E.L. (1992). Arguments and Cases: An Inevitable Intertwining. *Artificial Intelligence and Law*, 1(1):3–44.

[Stader, 1992] Stader, J. (1992). Applying neural networks. Technical Report AIAI-IR-11, AIAI, University of Edinburgh.

[Steinbart, 1985] Steinbart, P. (1985). *The Construction of an Expert System to Make Materiality Judgments*. Unpublished Ph.D. thesis, Michigan State University.

[Sviokla, 1990] Sviokla, J. (1990). ExperTAX: Coopers and Lybrand's Tax Accrual and Tax Planning Expert System. Technical report, Harvard Business School Case.

[Szolovits, 1983] Szolovits, P. (1983). Artificial intelligence in medicine. In *AAAS Selected Symposium*. Washington D.C.

[Tenor, 1988] Tenor, W. (1988). Expert systems for computer security. *Expert Systems Review*, 1(2):3–6.

[Trewin, 1991] Trewin, J. (1991). *The Impact of an Expert System on a Professional Accounting Organization*. Unpublished Ph.D. thesis, Michigan State University.

[Turbin & Watkins, 1986] Turbin, E. and Watkins, P. (1986). Integrating Expert Systems and Decision Support Systems. *Management Information Systems Quarterly*.

[Watkins & O'Leary, 1993] Watkins, P. and O'Leary, D. (1993). Integrating Intelligent Technologies into Conventional Information Systems. In *Expert Systems in Business and Finance: Issues and Applications*. John Wiley & Sons.

[White, 1988] White, H. (1988). Economic prediction using neural networks: the case of IBM daily stock returns. In *Proceedings of the IEE International Conference on Neural Networks*, volume 2, pages 451–459.

[Wielinga *et al*, 1992] Wielinga, B.J., Schreiber, A.Th. and Breuker, J.A. (March 1992). KADS: A Modelling Approach to Knowledge Engineering. *Knowledge Acquisition*, 4(1).

[Williamson, 1975] Williamson, O. (1975). *Markets and Hierarchies: Analysis and Antitrust Implications*. The Free Press, London.

[Willingham & Ribar, 1988] Willingham, J. and Ribar, G. (1988). *Development of an Expert Audit System for Loan Loss Evaluation*. Arthur Young, Reston, VA.

[Winston, 1984] Winston, P. (1984). *Artificial Intelligence*. Addison-Wesley, Reading, Ma.