Book Reviews

Suzanne Daningburg, Editor

Expert Systems, by Paul Harmon and David King. John Wiley and Sons, Inc., New York, NY, 1985. 283 pages, $21.00 (Canadian).

Reviewed by Kimiz Dalkir

While the emerging technology of artificial intelligence is of general interest to educational technology, this particular book is especially relevant as it is one of the few to make an explicit link between the two fields. The senior author, an instructional designer turned management consultant, provides a different perspective on the often Computer Science dominated literature on artificial intelligence. I highly recommend this book to anyone wishing an introduction to expert systems, or to anyone who wonders how Educational Technology can fit in with and best benefit from this new field.

Expert systems have enormous potential applications in education. However, these are usually only briefly alluded to in most books, and usually only within the restricted context of computer-aided instruction. Despite the subtitle, Artificial Intelligence in Business, which implies business applications, the scope is really quite diverse. As effectively stated in the forward:

> Everyone in business, from trainer to chief executive, must daily face problem solving and decision making based on extensive but incomplete, uncertain and even contradictory data and knowledge.

In short, there is something for everyone in this book and educational technologists can read through with a view towards adding expert systems to their already substantial repertoire of tools. One can also consider the field of artificial intelligence as an extension of the potential field of practice of Educational Technology. The authors point out that, given the shortage of people formally trained in the field, it is best to hire individuals with similar experience but in more traditional jobs. The three specific groups identified are: project managers, technical writers and instructional designers.

A broad overview of artificial intelligence, its (brief) history and potential applications are included in the first section, with particular emphasis on productivity and training issues. A good distinction is made between expert systems and knowledge-based systems, a boundary that is often blurred in the literature. An expert system is an embodiment of one or many experts' knowledge and expertise, in a particular domain. As such, it requires minimal expertise on the part of the user. The expert system is a mechanized consultant for problem solving. A knowledge-based system, on the other-hand, has much less inherent expertise.Thus, more know-how is required on the part of the user. In fact, such systems can often be used as a reference tool by an expert.

Training is identified as a good field for expert systems applications, since both training and expert systems require a high degree of expertise confined to a limited area of knowledge and/or skill. Expert systems can be used to decrease the amount of on-the-job-supervision, and consultations with manuals and with experts. The authors are less successful with their contention that expert systems are free of human biases, however. An expert system can be just as biased as any human expert because it is human knowledge and expertise that is modelled in the system. The old maxim "garbage-in-garbage-out" is equally valid whether applied to conventional programs or expert systems.

This selection does a good job of outlining the basic concepts and techniques of artificial intelligence and of expert systems, referring to MYCIN, an expert system to diagnose bacterial infections, as the standard system. This section goes beyond the standard introductions to expert systems, however, by tackling some of the thornier cognitive issues involved. There is a chapter on human problem solving, which addresses the question of where expertise comes into play and how analogous human and machine information processing really are. The complexities of expert systems are directly related to the complexities of problem solving strategies and cognitive functions in humans.

The authors delve further into the cognitive foundations of artificial intelligence by proposing a quantitative model of expertise, along with a very useful rule of thumb to ascertain whether or not expertise exists for a given area. This model incorporates a number of intriguing ideas, including a ten-year minimum requirement for building up expertise (at a rate of seven seconds per chunk) and a 70 millisecond period of time required to access any chunk of information once it is in long-term memory. These points, together with the more widely known limit of 5 to 7 chunks in short-term memory at any one time, are used to suggest that while there is no apparent limit on the amount of information that can be stored, there are limits on accessing this information (or remembering). Thus expertise consists not only of knowledge, but of knowledge access methods.

Using this premise, then, different knowledge representation and inference methodologies can be seen in a different light. A knowledge representation is the structuring and internal representation of all of the facts and rules of the application domain in an expert system. The inference strategy is how the system makes use of this knowledge base in order to answer queries and come up with recommendations. One interesting train of thought is to compare the search strategies used to solve problems or answer questions in human learners. For example, depth-first search will follow a promising branch of the knowledge tree to its end, then start over with another if the first was unsatisfactory, while breadth-first searching will try out several branches first and then follow the best one to its end. This may be analogous to learners who prefer to have an overall view of a subject matter first in contrast to individuals who prefer the building block approach to learning.

A similar analysis can be carried out for the two types of knowledge that can be represented in expert systems: deep and surface. Deep knowledge consists of first principles, axioms, facts and so on. This is analogous to school or textbook knowledge and is to be found in novices or experts confronted with a new type of problem. Surface knowledge is "compiled" in that it consists of heuristics or rules of thumb, judgment, intuition and so on. This is the core of expertise and is often acquired from a mentor or through experience.

The second section discusses expert system languages, tools and systems that are commercially available or under development. There is a good multi-dimensional taxonomy of artificial languages and tools, in addition to a useful checklist to evaluate such tools. Approximately fifteen tools are discussed in some detail, including an overview, the

knowledge representation and inference strategy used, the hardware required to implement it and the purchase price. This provides the reader with a good idea as to what is available and how to pick and choose among them.

Section 3 is a step-by-step guide to designing a small expert system, and more general guidelines on how to go about designing larger systems. A training media advisor is used as an example of a small system, with the complete program provided in M1 (an expert system development package). This example involves identifying the needs of trainees, simulating the job performance, identifying instructional factors to be considered and the budget constraints to be respected, in order to recommend the most cost-effective media to use for a particular training task. This section will thus be of particular interest to instructional designers.

Section 4 goes on to provide an overall view of the expert system market: where will expert systems be needed, by whom, and so on. The commercialization prospects of artificial intelligence and expected trends for the next five years are discussed. This section is geared more toward the entrepreneur who is interested in the business of developing and marketing expert system tools and applications, as well as consultants, who may wish to identify areas where their services are likely to be required.

The final section contains an excellent chapter on expert systems in training, which also deals with a number of issues that are are fundamental to educational technology. Among these are the variety of knowledge types that exist and their relationship to instructional strategies. A spectrum is proposed, with education (conceptual level) at one end, performance aids (specific task level) at the other, and training (procedural level) roughly in the middle. This is a good treatment of the often vague terms used in education, including "education" and "training." The authors identify three approaches to teaching, based on this spectrum:

1) to provide conceptual principles that allow the learner to think in abstract terms — *education* (e.g., sending a promising manager back to school to obtain an MBA;
2) to provide some theoretical information, but only within the context of carrying out a particular specified goal - *training* (e.g., sending an employee to a workshop or seminar to acquire particular techniques); and
3) to provide actual tools and a minimal amount of information to allow an employee to carry out a particular task — *job aid* (e.g., giving someone a computer and spreadsheet package to generate the *right* answer.

The authors thus make a distinction between education and training based on both the broadness of the instructional domain, with education having a more global realm than training which has a limited subset of this subject area, and based on the nature of the output (i.e., the individual who is taught). The trained person is much easier to identify than the educated one, which poses problems in defining the objective of education. If someone is successfully trained, then he or she is now capable of carrying out a particular job that they were not able to do before.

Some examples of specific training expert systems developed to date are also described in this section, including:

1) STEAMER - computer-assisted instruction to teach Naval officers how to run a ship's steam propulsion plant;
2) DEBUGGY - to identify student errors in solving math problems;
3) GUIDON - a tutorial on diagnosing bacterial infections (MYCIN rearranged);
4) DELTA - assists mechanics in diagnosing and repairing train engines; and
5) PUFF - a pulmonary diagnosis instructional system.

In general, I found this book to be quite interesting to read, with many excellent diagrams, figures, tables and graphs to complement the points raised in the text. The overview of expert systems is a good introduction to the subject, although a more balanced view, pointing out the limitations of expert systems and common pitfalls in their development and use, would have been preferable. The reader should thus maintain a healthy degree of skepticism regarding any blanket statements extolling the virtues of these systems.

There are three appendices which may also prove quite useful. The first is a glossary of the terms prevalent in any artificial intelligence publication. Another is a list of companies involved in expert system research, along with their addresses and contact persons to write to if the reader is interested in further information. The annotated bibliography at the end of is also a good starting point for anyone interested in pursuing the subject further, as I hope many will. It is likely that if any synergistic association is to formed between artificial intelligence and educational technology, it will be at the initiative of the latter, instead of the leisure of the former.

Kimiz Dalkir is a Knowledge Engineer (acquisition and representation of expert knowledge) at the Canadian Workplace Automation Research Centre (funded by the DOC) in Montreal, Quebec and a doctoral student in Educational Technology at Concordia University.