Potential Applications of Knowledge-Based Expert Systems in Transportation Planning and Engineering

CHE-I YEH, STEPHEN G. RITCHIE, and JERRY B. SCHNEIDER

ABSTRACT

The objectives of this paper are to describe the characteristics of knowledge-based expert systems (KBES) and to suggest some applications that appear to have a high potential for development in the field of transportation planning and engineering. Such systems represent a rapidly developing branch of artificial intelligence (AI) and computer science that is already having significant impacts in many disciplines. KBES use interactive computer programs that seek to provide a level of performance and expertise that is matched by only a few human experts in a particular problem domain. An overview is provided of AI and KBES concepts and of existing KBES and their architecture. The current scope of expert systems is described in an attempt to identify high-potential applications in the fields of transportation planning and engineering. A number of these applications are identified and discussed. It is concluded that the potential appears high for KBES to become useful tools for practicing transportation planners and engineers.

Since the mid-1970s one of the most significant accomplishments in the field of artificial intelligence (AI) has been the development of knowledge-based expert systems (KBES). These systems are interactive computer programs that employ a collection of judgment, experience, rules-of-thumb, intuition, and other expertise in a particular field, coupled with inferential methods of applying this knowledge, to provide expert advice on the performance of a variety of tasks (1,2). Domain-independent reasoning techniques were given concentrated attention in the early history of AI. The General Problem Solver (GPS) (3), a classic example, could solve many puzzles and prove theorems. However, it was soon recognized that such universal AI programs could not handle very many real-world problems, so AI researchers have shifted their attention to the construction of expert systems with domain-dependent knowledge (4).

Many KBES are problem-solving programs designed to reach the level of performance of a human expert in a specific professional domain. Operational expert systems have already been developed in a number of disciplines. These include MYCIN for medical consultation (5), DIPMETER for oil well logging (6), PROSPECTOR for mineral exploration (7), and R1 for computer system configuration (8). The utility of these new computer programs has been proven in their short histories. For example, PROSPECTOR discovered a rich lode of molybdenum ore in eastern Washington State. Although this area had been explored by many experts since World War I, none had ever found the exact site of the ore. In the early 1980s PROSPECTOR was used to make inferences from data in a partly explored area and predicted ore in a particular location. Subsequent drilling confirmed predictions of where ore would and would not be found. The reason that PROSPECTOR could solve this difficult problem was that it not only incorporated the knowledge acquired from nine expert geologists but also had the advantages of the computer--high-speed operation and huge memory capabilities--that it employs during its inferencing activities. Human experts find it difficult to deal with more than four or five data elements at the same time although their knowledge in this field may be as good as that of PROSPECTOR (9).

KBES have high potential for solving problems that lack explicit algorithms (e.g., problems for which a numerical model does not exist). Currently, ill-defined problems can only be solved in a limited fashion by employing an algorithmic approach after extensive knowledge is summarized and a simple model created. However, there are generally some experts in every field who can solve such ill-structured problems using their past experience and knowledge. Many important real-world problems are ill-structured (e.g., designing an optimal transit route structure or making decisions about how to rehabilitate a major highway). Explicit and well-formulated algorithms do not exist in these domains. The transportation field, in particular, is full of such ill-structured problems in which human behavior, social and political considerations, and multiobjective decision making are involved.

In this paper an overview of expert systems is presented and some potential applications in the field of transportation are suggested. First, the basic concepts and structure of expert systems are introduced. The types and scope of expert systems are described. Criteria that can be used to ascertain when to apply expert systems are given, and the tasks and tools involved in building an expert system are discussed. After this overview of expert systems, some high-potential applications in transportation planning and engineering are identified and discussed.

C.-I Yeh, Institute of Transportation Studies, University of California, Irvine, Calif. 92717. S.G. Ritchie, Department of Civil Engineering and Institute of Transportation Studies, University of California, Irvine, Calif. 92717. J.B. Schneider, Department of Civil Engineering, University of Washington, Seattle, Wash. 98195.

TYPES AND SCOPE OF EXPERT SYSTEMS

Types

On the basis of the different functions they perform, expert systems can be divided into several types (10):

1. Diagnosis and debugging: inferring malfunctions from observed information;
2. Design: developing and configuring a process for a set of specific objectives that satisfy given constraints;
3. Planning: identifying a set of actions that will lead to the achievement of given objectives in the future;
4. Interpretation: explaining the characteristics of input and output data;
5. Forecasting: inferring likely future outcomes given a set of past trends;
6. Monitoring: observing and comparing observations with planned and desired system characteristics;
7. Repair: developing a plan to recover from a failure or malfunction; and
8. Control: a most complicated system that interprets the data, diagnoses the problem, predicts the future, formulates a plan, executes the plan, and monitors its implementation.

Many actual transportation problems are quite broad and involve aspects of several of these problem types. For example, typical transportation problems often involve the need to interpret incomplete and faulty data, predict ridership or population, diagnose traffic congestion, design new routes or reconstruct an existing network, monitor roadway traffic, and control network operations.

Scope

Not all problems are suitable candidates for the expert systems approach. A list of criteria that can be used to evaluate potential expert systems applications follows.

1. Algorithmic solutions are impractical because of complex physical, social, political, or judgmental components, which generally resist precise description and deterministic analysis.
2. Faulty or incomplete data will be encountered during the problem-solving activity.
3. Experts need technical tools (such as handbooks and computers) and reference resources to identify, make inferences, and analyze the problem because of the extensive basic and background knowledge requirements.
4. The problem-solving model must be changed to handle different types of problems or to adapt to dynamic conditions.
5. High-performance results are required in a short time, whereas a complete numerical solution could be obtained only after spending unreasonable or unacceptable amounts of time and money using traditional methods, and a reliable "short-cut" traditional method does not exist.
6. A failure by a faulty human expert cannot be tolerated.
7. Knowledge transfer from scarce human experts is too difficult or costly or may take too long.
8. A human expert is not available.
9. Public knowledge (books, papers) is not sufficient to allow most persons to reach an expert's level. Substantial private knowledge (expertise, heuristics) exists only in the minds of the experts.
10. The potential payoffs are high.

BASIC CONCEPTS

Expert Systems Versus Traditional Computer Programs

Most traditional computer programs involve the use of algorithmic procedures. These algorithmic programs can only follow step-by-step commands or procedures in searching for a solution to a problem. An existing model, usually a numerical model, is needed to support the operation of the search algorithm. The developer of the program must predetermine the sequence of problem-solving procedures and ensure that the set of commands can handle all possible combinations of problems or subproblems to obtain a unique and correct result. However, the knowledge pertinent to the problem and the process used to employ this knowledge are often interrelated (2). This produces some of the difficulties encountered with traditional programs:

1. An algorithm for solving the problem must exist. Many problems are impossible or impractical to deal with because the knowledge is difficult to represent in a numerical form and a sequence of steps that will produce a solution is unknown.
2. Only problems for which correct and complete data are available can be handled, otherwise an incorrect result will be obtained.
3. The maintenance and updating of a traditional program are difficult because the knowledge is dispersed throughout the entire program and intertwined with the algorithmic computation process.

Expert systems are designed to overcome these difficulties of traditional programs. A separate knowledge processor determines when, how, and where to apply every individual element of knowledge. An explicit problem-solving algorithm is not needed. Most expert systems can handle incomplete or faulty data by making inferences from their own knowledge base. Maintenance of knowledge is also more easily accomplished for expert systems because the knowledge base is separated from the control process, as discussed next.

Architecture of Expert Systems

In an expert system the program is divided into a knowledge base and a control process for applying and selecting knowledge (also known as an inference engine or inference machine). The general architecture of KBES is shown in Figure 1. Each component of the system (10-13) is described.

Knowledge Base

The knowledge base represents the power of the expert system. The knowledge base contains all relevant information, facts, or causal knowledge of the domain, and rules used for problem-solving activities to determine what actions will be executed if certain situations are met. There are various techniques for representing knowledge. The most widely used in current expert systems are rules, frames, and semantic nets (13). The most common method, particularly in transportation expert systems developed to date, involves production rules. These have an IF condition THEN action format. When the IF portion or premise of a rule is satisfied by the facts, the action specified by the THEN portion is performed. The rule is then said to "fire." For example, one rule in SCEPTRE (14), an expert system for pavement rehabilitation decision making, states:

IF Depth of rutting is greater than 3/4 in. and
Transverse cracking is the predominant type of cracking

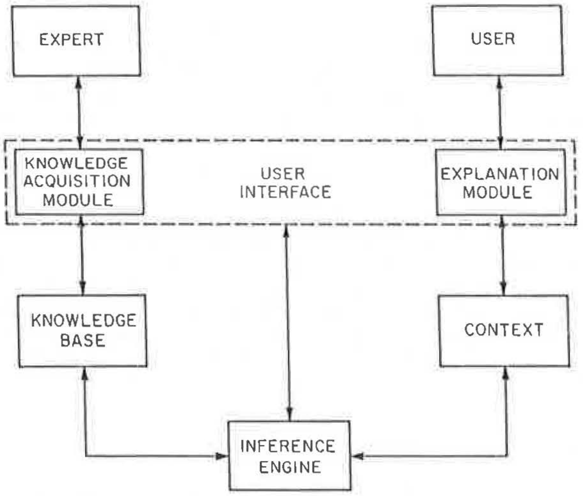

FIGURE 1 Architecture of a knowledge-based expert system.

THEN Mill and replace
or Prelevel and thin overlay
or Prelevel and medium overlay

Context

The context is the data level of an expert system. It contains all data, symbols, or facts that reflect the current status of the solution. This information may describe the problem being solved, the rules that are employed, and the facts that are true. The context is also called short-term memory.

Inference Machine

The inference machine (also called inference engine) is the control level of an expert system. The objective of the inference machine is to find a conclusion for a subgoal or the answer to an entire problem. It searches for facts through the knowledge base and identifies new facts for subsequent inferencing.

Explanation Module

Unlike a traditional program, which is often a "black box" to the user, an expert system has an explanation module that explains its problem-solving strategy to the user.

Knowledge Acquisition Module

Knowledge acquisition is often the bottleneck of the entire expert system development process. An automatic knowledge acquisition module can speed up the development of KBES. However, a system-independent knowledge acquisition module does not currently exist because each problem domain has its own knowledge structure.

User Interface

The user accesses the system through a user interface, which should be friendly so that man and machine can communicate directly and efficiently. Currently, actual language processors do not exist, so most user interfaces are implemented as a problem-oriented subset of English.

Building an Expert System and Associated Tools

The major task of building an expert system is to transfer the expertise and knowledge acquired from one or more experts to a computer program. The mission of expert systems developers (knowledge engineers) is to carry out such a transformation and to ensure that the performance of the resulting expert system can reach the desired level. The steps involved in constructing an expert system will vary depending on the characteristics of the problem, the objectives selected, and the development tools available. However, the following stages are normally encountered in the development of an expert system (10):

1. Identification: The first step in building an expert system is to identify the area, concepts, and characteristics of the problem and solution. In addition, the participants and resources (time, labor, and computing facilities) needed during development should be identified.
2. Conceptualization: The concepts needed to represent knowledge and the overall structure of knowledge-control strategies must be determined before a preliminary system design can be completed.
3. Formalization: This stage involves design of the formal organization of the knowledge consistent with the development tools or languages used. The detailed design of the system is formulated by the expert or experts and the knowledge engineer or engineers.
4. Implementation: In this stage the knowledge engineer turns the formalized knowledge into a working computer program.
5. Testing: The performance and behavior of the prototype expert system are iteratively evaluated through comparison with the human expert's abilities. Revisions of the system are then made by a knowledge engineer using additional advice from the human expert or experts.

In building expert systems, symbol-manipulation programming languages such as LISP and PROLOG have been widely used. These languages have been specially designed for artificial intelligence applications. However, a much faster route to system development can be to use a knowledge engineering toolkit or "shell," which comprises an inference engine, empty knowledge base and context structure, and support facilities such as a knowledge-base editor and user explanation facility. The system developer must then enter the rules (in a rule-based system) in the knowledge base. In recent years, a large number of commercial shells have become available for microcomputers. Many of these shells are also being coded in conventional programming languages such as C or PASCAL (for improved expert system execution speed), although knowledge-base creation and editing may still involve an English-language programming style.

POTENTIAL APPLICATIONS OF EXPERT SYSTEMS IN TRANSPORTATION

As discussed previously, KBES seek to provide a level of performance that is similar to that of an expert, or group of experts, in a particular problem domain. Such systems are primarily applicable to situations that require specialized knowledge, skill, experience, or judgment for determination of a solution or development of a solution strategy. In such cases, the problem is usually said to be ill-structured, in the sense that a numerical algorithmic solution is not available or is impractical. Because so many of the problems that transportation professionals face are of this kind, it can be said that, in general, the potential appears high for KBES to become useful tools for the practicing transportation planner and engineer. KBES can be envisaged as functioning as expert consultants, capable of explaining their reasoning and why they arrived at certain conclusions. Eventually, users may be able to learn from an expert system in the same way people can now learn from a dialogue with an expert consultant.

Recent applications of expert systems in civil engineering are described elsewhere (15), although none of the applications presented involve primarily transportation examples. Considerable research is now under way on developing expert systems in the transportation field, but comparatively little work has been reported to date. Some initial efforts include DELTA, a fault diagnosis expert system for troubleshooting diesel-electric locomotive maintenance problems (16); an automatic crew-scheduling expert system used by the National Aeronautic and Space Administration (NASA) for the space shuttle (2); DIRECTOR, an intelligent front-end expert system for an educational urban transportation decision-making simulation model; and identification of potential applications of expert systems in air traffic control (17). Recent operational prototype expert systems in transportation include SCEPTRE, for pavement rehabilitation decision making (18); CHINA, for acoustic design of highway noise barriers (19); TRALI, for traffic signal setting assistance (20); HERCULES, for postdisaster traffic recovery strategies (21); and EXPERT-UFOS, for multicriteria traffic network design (22).

Although a number of criteria for evaluating potential expert system applications have been outlined, it is still a difficult task to forecast developments and applications (without the aid of an expert system). Nevertheless, in this section of the paper, an attempt is made to identify a number of high-potential applications of KBES in the field of transportation planning and engineering.

Design

Facility Design

As in many other design problems, design of transportation facilities involves application of considerable expertise, experience, and heuristic expert knowledge. It is clearly a high-potential application area for an expert system. For example, a roadway design system could not only reduce design costs but could also provide a high-quality professional design service. Such a complete expert system would need to employ more complex AI techniques than are currently available. However, some subproblems could be examined using expert systems techniques to gradually develop the necessary overall capabilities.

In addition, well-designed intersections can enhance safety and reduce traffic delay and congestion. An expert system that could infer traffic demands and relate physical, environmental, and other design constraints to aid in the design of a high-quality intersection would be of substantial use to many smaller cities that cannot afford to hire expensive expert consultants.

An existing expert system in this general category is CHINA (19). CHINA addresses the problem of acoustically designing a highway noise barrier. It contains the expert knowledge of several specialists in the control of highway noise and can act as an expert advisor to the novice engineer or as a colleague to more experienced engineers on complex abatement problems. TRALI is another example system, which provides assistance to traffic engineers designing traffic signal settings for isolated signalized intersections (20). This is a classic and common problem in transportation engineering. TRALI addresses a particular shortcoming of existing design aids that cannot deal with uncommon geometries. EXPERT-UFOS is an expert system that aids in designing optimal capacity additions and deletions from a transportation network (22). The system addresses network design problems that must be evaluated using multiple conflicting criteria. EXPERT-UFOS achieves this by incorporating a multicriteria evaluation method in the expert system.

Study Design

Data collection and sampling are needed to solve many transportation problems. A well-designed study and data collection program can provide a controlled environment and inference source for transportation planners and engineers. However, only a limited number of experienced engineers and planners can perform such activities well. An expert system that could make inferences about data sources and data types to suggest proper data collection methodologies and sampling techniques could substantially improve the quality and efficiency of local study designs.

Planning

Methodology

The choice of a methodological approach for urban transportation planning substantially affects the scope, duration, data intensiveness, policy sensitivity, and cost of the planning effort. The judgment and experience required to select appropriate analysis tools, such as system, sketch, or quick-response planning approaches, suggests a useful role for a planning methodology advisor. Also, the initial generation of transportation alternatives for subsequent analysis and evaluation is a critical but not well-defined step. An expert system could assist planners by relating appropriate transportation technologies and system management and control options to the intensity and distribution of trips, land uses, and socioeconomic characteristics in the study area.

Disaster Response Planning

Disaster response planning is a concern of all levels of government. To reduce the effects of disaster, it is often necessary to evacuate people from the disaster area and to send assistance teams and supplies into the area as soon as possible. Such a response depends on finding ways to use the transportation system effectively when it is damaged, badly congested, and not likely to be repaired quickly. An expert system that incorporates the knowledge of disaster response experts and local conditions and resources could be of substantial assistance in dealing with the dynamics of such situations.

Bus Transit Network Planning

Current urban transit bus networks are full of inefficiencies that are difficult to identify and even more difficult to remove without creating worse problems in other parts of the network. Yet some transit specialists can examine current route and schedule layouts and come up with diagnoses that lead to useful definitions of problems and corrective measures. If the capabilities of these persons could be incorporated in an expert system to aid local planners, productivity improvements in transit system performance could be expected.

Operations and Control

Traffic Congestion Diagnosis

The causes of traffic congestion are highly dependent on a variety of physical, environmental, operational, geometric, land use, human, and many other factors. Some experts can determine the major reasons for traffic congestion from their observations and knowledge. However, these experts are scarce and expensive. A traffic congestion diagnosis expert system that incorporates the knowledge and judgmental abilities of several experts could diagnose traffic congestion just like a human expert. The performance of the expert system would depend heavily on the expertise acquired from the specific experts and the manner in which it was integrated and organized.

The expert system HERCULES (21) addresses postdisaster traffic recovery and generates traffic control plans that would make good use of the links remaining in a damaged urban road network. HERCULES attempts to keep postdisaster congestion at tolerable levels by recommending control plans that limit allowable volumes on specific links.

Roadway Safety Diagnosis

The safety of a roadway is an important issue for transportation engineers. Many accidents are caused by poor design. There is a possibility that such engineering mistakes underlie many accidents. It is not possible, in practice, for every accident to be reviewed by a safety expert. An expert system that can detect any engineering design factors that may have contributed to an accident could provide a useful and inexpensive roadway safety diagnosis capability.

Hazardous Material Transportation

The problem of hazardous material transportation involves complex and multiple-objective decision making in a sociopolitical context that is difficult to represent in a simple numerical model. An expert system could provide assistance to such decision making and would be based on an extensive knowledge base of hazardous materials properties, population density and distribution, and roadway characteristics.

Air Traffic Control

The human air traffic controller currently plays a major role in the air traffic system. Handling increasing traffic volume safety places controllers under great pressures and much stress. Any mistake that the air controller makes could obviously cause many deaths. An air traffic expert system that incorporated appropriate heuristics, perhaps combined with algorithmic procedures, could provide much assistance to the human controllers. However, such an expert system must be tested completely before being placed in service. As long as the knowledge built into the expert system is comprehensive and organized correctly, the probability of a system failure or mistake can be reduced to a minimum. An air traffic control expert system can monitor human controller behavior and function as a warning or backup system to reduce the pressure on the human controller. Work on this topic is currently under way (17).

Ground Traffic Signal Timing Control

As expert system that could provide expertise on the timing and coordinating of traffic signals would be helpful for many smaller cities. Such a system could provide a professional opinion to help local traffic engineers do their jobs more effectively. The characteristics of ground traffic control are in some ways more complex than those of air traffic control. Unlike air traffic control, which employs radar to monitor all elements of the system, the ground traffic control problem is poorly defined. Parts of the problem can be defined by the use of cameras, sensors, traffic controllers, and observers, but the data are often incomplete and may be faulty or conflicting. Such characteristics relate well to the capabilities of expert systems. An expert system that performs like a human expert at an isolated signal set can reduce manpower needs and provide continuous day and night service. Most of all, such an expert system would provide the best quality of service if the knowledge base were acquired from the best experts and modified to fit local conditions. Such an expert system has a different approach than do current algorithmic control systems, which usually cannot handle special events, emergency vehicle needs, unusual geometries, and risky situations. Another major difference is that algorithmic systems cannot handle incomplete and faulty data. It appears likely that a ground traffic control expert system that can control all of the traffic signals in a designated area would be most helpful in obtaining higher performance in the future.

Dispatching and Scheduling

It is clear that some dispatchers and schedulers can perform these tasks much better than most others. These people are experts at their jobs. It is normally difficult to transfer such expertise to other people or other locations. Usually, the knowledge that makes these experts better than their peers is difficult to represent in an explicit algorithm. An expert system that contains their knowledge could assist in solving the problem of knowledge transfer and improve local capabilities.

Maintenance and Rehabilitation

Transit Vehicle Maintenance

Many transportation systems use complex mechanical equipment. The sudden failure of such machines may cause critical problems. For example, the mechanical failure of an automated guideway transit system (people mover) may cause delays or system shutdown and, possibly, accidents. In most situations, faults must be repaired or a backup system invoked in the shortest possible time. Human experts may not be able to respond fast enough. An expert system that contains knowledge of all repair tasks and backup capabilities could be a valuable tool for transportation providers that cannot afford to have frequent failures on and off the system.

Pavement Rehabilitation

In recent years, continued deterioration of the highway infrastructure has led to increased emphasis on pavement rehabilitation, with national annual expenditures in the United States of billions of dollars. Given the pressure of limited budgets and the need to address extensive decay problems, decisions on pavement rehabilitation have become critical. There are only a few experts who, using field observations and their own knowledge and experience, can determine the causes of pavement decay and appropriate solution strategies with high confidence. Because the analysis and design of project-specific rehabilitation strategies rely so heavily on expert pavement engineers, and the tasks involved are both complex and ill-defined, conventional computer tools are of limited use.

SCEPTRE (18) is an expert system for flexible pavement rehabilitation decision making. Working from user inputs and a knowledge base constructed from several human experts, the system can deduce a set of feasible project-level rehabilitation strategies for subsequent detailed analysis and design. The system can also readily explain its reasoning and conclusions and is easily modified. It can therefore make its body of specialized knowledge accessible to a broad range of potential engineering users.

CONCLUDING COMMENTS

The objectives of this paper were to describe the characteristics of knowledge-based expert systems and to identify several high-potential applications in the field of transportation planning and engineering. As discussed earlier in the paper, not all problems are suitable candidates for the expert systems approach. A number of criteria were proposed for evaluating potential expert systems applications, and the steps involved in constructing an expert system were discussed. The most difficult phase of this process remains the acquisition and transformation of knowledge and expertise from one or more experts to the knowledge base. Although relatively little work on KBES in the transportation area has been reported to date, research is under way and can be expected to grow. In several other disciplines, operational expert systems have already been developed and their utility proven.

Because so many of the problems transportation professionals face require specialized knowledge, skill, experience, and judgment for determination of solution strategies, the authors believe that, in general, the potential appears high for KBES to become useful tools for practicing transportation planners and engineers. A number of potential applications were identified in the areas of design, planning, operations and control, and maintenance and rehabilitation. However, given the breadth of the transportation field, there may be many domains not included in this paper that also represent high-potential applications. To identify new applications and research needs, consultations with appropriate experts and a more careful and complete review of domain-dependent problems are required. Such work should be followed by development and evaluation of new prototype expert systems. This would improve the ability to assess the feasibility and true potential of such systems in transportation planning and engineering.

ACKNOWLEDGMENT

This research was partly supported by the National Science Foundation.

REFERENCES

1. J. Gasching, R. Reboh, and J. Reiter. Development of a Knowledged-Based Expert System for Water Resource Problems. SRI International, Menlo Park, Calif., Aug. 1981.
2. W.B. Gevarter. Expert Systems: Limited But Powerful. Spectrum, IEEE, Aug. 1983, pp. 39-45.
3. A. Newell and H.A. Simon. Human Problem Solving. Prentice-Hall, Inc., Englewood Cliffs, N.J. 1972.
4. M. Konopasek and S. Jayaraman. Expert Systems for Personal Computers. BYTE, Vol. 9, No. 5, May 1984, pp. 137-154.
5. E.H. Shortliffe. Computer-Based Medical Consultation: MYCIN. American Elsevier, N.Y., 1976.
6. R. Davis et al. The Dipmeter Advisor: Interpretation of Geological Signals. Proc., Seventh International Joint Conference on Artificial Intelligence, Vancouver, British Columbia, Canada, 1981.
7. R.O. Duda et al. A Computer-Based Consultant for Mineral Exploration. SRI International, Menlo Park, Calif., 1979.
8. J. McDermott. Rl: A Rule-Based Configurer of Computer Systems. Technical Report CMU-CS-80-119. Department of Computer Science, Carnegie-Mellon University, Pittsburgh, Pa. 1980.
9. Artificial Intelligence is Here. Business Week, July 9, 1984.
10. F. Hayes-Roth, D.A. Waterman, and D.B. Lenat, eds. Building Expert Systems. Addison-Wesley, Reading, Mass., 1983.
11. D.R. Rehak. Expert Systems in Water Resource Management. Proc., ASCE Conference on Emerging Techniques in Storm Flood Management, Niagara-on-the-Lake, Ontario, Canada, 1983.
12. D. Sriram et al. Expert Systems in Civil Engineering--A Survey. Report R-82-137. Department of Civil Engineering, Carnegie-Mellon University, Pittsburgh, Pa., July 1982.
13. D.A. Waterman. A Guide to Expert Systems. Addison-Wesley, Reading, Mass., 1986.
14. S.G. Ritchie, C. Yeh, J.P. Mahoney, and N.C. Jackson. Development of an Expert System for Pavement Rehabilitation Decision-Making. In Transportation Research Record 1070, TRB, National Research Council, Washington, D.C., 1986, pp. 96-103.
15. C.W. Kostem and M.L. Maher, eds. Expert Systems in Civil Engineering. Symposium Proceedings, American Society of Civil Engineers, Seattle, Wash., 1986.
16. P.P. Bonissone and H.E. Johnson. Expert System for Diesel Electric Locomotive Repair. Knowledge-Based Systems Report. General Electric Co., Schenectady, N.Y., 1983.
17. G.D. Gosling. Application of Expert Systems in Air Traffic Control. Submitted to ASCE Journal of Transportation Engineering, 1986.
18. S.G. Ritchie. Expert Systems in Pavement Management. Transportation Research, 1987 (in press).
19. R.A. Harris, L.F. Cohn, and W. Bowlby. An Application of Artificial Intelligence in Highway Noise Analysis. In Transportation Research Rec-

ord 1033, TRB, National Research Council, Washington, D.C., 1985, pp. 11-15.
20. C. Zozaya-Gorostiza and C. Hendrickson. An Expert System for Traffic Signal Setting Assistance. Journal of the Transportation Engineering Division, ASCE, 1986 (in press).
21. C. Yeh. Development of a Post-Disaster Transportation Network Trip Simulation Model and Traffic Recovery Strategy Expert System. Ph.D. dissertation. University of Washington, Seattle, 1985.
22. S. Tung. Designing Optimal Networks: A Knowledge-Based Computer-Aided Multicriteria Approach. Ph.D. dissertation. University of Washington, Seattle, 1986.

Publication of this paper sponsored by Committee on Computer Graphics and Interactive Computing.