NATIONAL BUREAU OF STANDARDS

NBS SPECIAL PUBLICATION 728

U.S. DEPARTMENT OF COMMERCE/National Bureau of Standards

# Expert Systems and Emergency Management: An Annotated Bibliography

S.I. Gass, S. Bhasker, and R.E. Chapman

FEDERAL EMERGENCY MANAGEMENT AGENCY

PACE AC BELLO MERITA

The National Bureau of Standards[1] was established by an act of Congress on March 3, 1901. The Bureau's overall goal is to strengthen and advance the nation's science and technology and facilitate their effective application for public benefit. To this end, the Bureau conducts research and provides: (1) a basis for the nation's physical measurement system, (2) scientific and technological services for industry and government, (3) a technical basis for equity in trade, and (4) technical services to promote public safety. The Bureau's technical work is performed by the National Measurement Laboratory, the National Engineering Laboratory, the Institute for Computer Sciences and Technology, and the Institute for Materials Science and Engineering.

## *The National Measurement Laboratory*

Provides the national system of physical and chemical measurement; coordinates the system with measurement systems of other nations and furnishes essential services leading to accurate and uniform physical and chemical measurement throughout the Nation's scientific community, industry, and commerce; provides advisory and research services to other Government agencies; conducts physical and chemical research; develops, produces, and distributes Standard Reference Materials; and provides calibration services. The Laboratory consists of the following centers:

- Basic Standards[2]
- Radiation Research
- Chemical Physics
- Analytical Chemistry

## *The National Engineering Laboratory*

Provides technology and technical services to the public and private sectors to address national needs and to solve national problems; conducts research in engineering and applied science in support of these efforts; builds and maintains competence in the necessary disciplines required to carry out this research and technical service; develops engineering data and measurement capabilities; provides engineering measurement traceability services; develops test methods and proposes engineering standards and code changes; develops and proposes new engineering practices; and develops and improves mechanisms to transfer results of its research to the ultimate user. The Laboratory consists of the following centers:

- Applied Mathematics
- Electronics and Electrical Engineering[2]
- Manufacturing Engineering
- Building Technology
- Fire Research
- Chemical Engineering[2]

## *The Institute for Computer Sciences and Technology*

Conducts research and provides scientific and technical services to aid Federal agencies in the selection, acquisition, application, and use of computer technology to improve effectiveness and economy in Government operations in accordance with Public Law 89-306 (40 U.S.C. 759), relevant Executive Orders, and other directives; carries out this mission by managing the Federal Information Processing Standards Program, developing Federal ADP standards guidelines, and managing Federal participation in ADP voluntary standardization activities; provides scientific and technological advisory services and assistance to Federal agencies; and provides the technical foundation for computer-related policies of the Federal Government. The Institute consists of the following centers:

- Programming Science and Technology
- Computer Systems Engineering

## *The Institute for Materials Science and Engineering*

Conducts research and provides measurements, data, standards, reference materials, quantitative understanding and other technical information fundamental to the processing, structure, properties and performance of materials; addresses the scientific basis for new advanced materials technologies; plans research around cross-country scientific themes such as nondestructive evaluation and phase diagram development; oversees Bureau-wide technical programs in nuclear reactor radiation research and nondestructive evaluation; and broadly disseminates generic technical information resulting from its programs. The Institute consists of the following Divisions:

- Ceramics
- Fracture and Deformation[3]
- Polymers
- Metallurgy
- Reactor Radiation

[1]Headquarters and Laboratories at Gaithersburg, MD, unless otherwise noted; mailing address Gaithersburg, MD 20899.
[2]Some divisions within the center are located at Boulder, CO 80303.
[3]Located at Boulder, CO, with some elements at Gaithersburg, MD.

# Expert Systems and Emergency Management: An Annotated Bibliography

Saul I. Gass
Suneel Bhasker
College of Business and Management
University of Maryland
College Park, Maryland 20742

and

Robert E. Chapman
Center for Applied Mathematics
National Engineering Laboratory
National Bureau of Standards
Gaithersburg, Maryland 20899

Sponsored by:

Federal Emergency Management Agency
500 C Street, S.W.
Washington, DC 20472

and

Center for Applied Mathematics
National Engineering Laboratory
National Bureau of Standards
Gaithersburg, Maryland 20899

U.S. DEPARTMENT OF COMMERCE, Malcolm Baldrige, Secretary
NATIONAL BUREAU OF STANDARDS, Ernest Ambler, Director

November 1986

Library of Congress Catalog Card Number: 86-600591

National Bureau of Standards Special Publication 728
Natl. Bur. Stand. (U.S.), Spec. Publ. 728, 179 pages (Nov. 1986)
CODEN: XNBSAV

U.S. Government Printing Office
Washington: 1986

For sale by the Superintendent of Documents, U.S. Government Printing Office, Washington, DC 20402

ABSTRACT

This report is the result of an in-depth review of the recent technical literature on expert systems. The material contained in this report provides a basis for assessing the potential for using expert systems in emergency management operations. In choosing the material for inclusion in this report, special emphasis was placed on those aspects of expert systems which addressed the types of problems encountered in emergency management operations.

The report is designed for use as a resource document and as a tutorial on expert systems and emergency management. Each chapter consists of a brief topic essay, followed by a set of references which expand on the main themes of the essay.

Keywords: Artificial intelligence; emergency management; expert systems.

Table of Contents

Table of Contents

## List of Tables



## List of Figures

# 1. INTRODUCTION

## 1.1 THE PURPOSE OF THIS REPORT

Improving the management of emergency situations at all levels of government has always been a matter of concern. Thus, emergency managers have been quick to search out and utilize new equipment and technologies. For example, our ability to respond to the ever changing requirements of an emergency has improved greatly due to the use of modern information and communications systems. Today, with the increasing use of personal and other types of computers, coupled with recent advances in artificial intelligence, the emergency manager is faced with the difficult task of determining what aspects of these advanced technologies should be investigated in terms of their role in the management of emergencies. Unlike past improvements, the blend of computers and artificial intelligence, especially the subfield of expert systems, is not well understood. Also, the development and testing of such computer-based systems can be quite costly; there is little experience to guide us. However, as the application of expert systems to emergency management operations could result in truly major advances, it is important for the community of emergency managers, administrators and researchers to become familiar with the potential of expert systems. This annotated bibiliography is an attempt to provide the necessary information with which one can review and further pursue the current status of expert systems, with emphasis on its relation to emergency management.

We have included source material that deals with: (1) the nature of expert systems; (2) available methodologies and tools for building expert systems; (3) the logic and reasoning behind the tools used in expert systems; (4) the approaches used to measure uncertainty of knowledge; (5) the issues of user interface; and (6) the current and potential emergency management uses of expert system methodologies. The references will provide the reader with different views of the theoretical and applications viability of expert systems. This applies especially to the concepts that form the core of current expert sytems techniques -- knowledge representation, logic and reasoning, and uncertainty. It is hoped that this bibliography will aid the reader to focus on those papers and books that will be of most value with respect to the reader's background and interests.

## 1.2 LITERATURE SURVEY

Although it was our aim to find and organize all material on expert systems and emergency management, this proved to be an impossible task due to the vast amount of relevant literature. The field of expert systems, in particular, was demonstrating exponential growth in terms of publications. Also, as we began gathering and reviewing books and papers, it became clear that much of the writings on expert systems were short-lived, represented experimental ventures, and became obsolete rather rapidly. Thus, we found it both convenient and appropriate to emphasize the inclusion of material published within the last five years. Care was taken, however, not to overlook seminal works that by definition have stood the test of time. Further, in many discussion areas, two or more items that differ only in depth or level of technical presentation were included so as to make the bibiliography more useful to the general emergency management community.

The actual literature search was done in two stages. First, the computer-based DIALOG Information Retrieval Service at the National Bureau of Standards was used to retrieve items from the following files:

- NTIS, which is an extensive information base consisting of governmental-sponsored research, development, and engineering reports plus analyses, journal articles, and translations prepared by Federal agencies, their contractors or grantees. This information is produced and maintained by the National Technical Information Service of the U.S. Department of Commerce. It is the means by which unclassified, unlimited distribution reports are made available to the public. The NTIS data base includes material from both the physical and social sciences, including topics of immediate widespread interest, such as environmental pollution and control, energy conservation, technology transfer, health planning, societal problems, and urban and regional development and planning. The file covers the period from 1960 to the present. The on-line reviews of NTIS produced nearly 300 references. The majority of the references produced by this search were theoretical and technical in scope.

- INSPEC is the largest English language information base in the fields of physics, electrotechnology, computers, and control. It corresponds to the printed Physics Abstracts, Electrical and Electronics Abstracts, and Computer and Control Abstracts. Non-English language source material is also included, but is indexed and abstracted in English. Journal papers, conference proceedings, technical reports, books, patents, and university theses are abstracted for inclusion in INSPEC. The total number of journals scanned is over 2000; 200 of these are abstracted completely. INSPEC covers citations from 1969 to the present. During the on-line review of INSPEC, nearly 200 references were produced. This search produced references that were more varied in the types of issues dealt with.

- The COMPENDEX data base is the machine-readable version of Engineering Index; it provides abstracted information from the world's significant engineering and technological literature. It has world-wide coverage of approximately 3500 journals, publications of engineering societies and organizations, papers from the proceedings of conferences, and selected government reports and books. COMPENDEX covers citations from 1970 to the present. More than 50 references were produced during the on-line review of COMPENDEX.

- The Ei ENGINEERING MEETINGS is a data base that covers significant published proceedings of engineering and technical conferences, symposia, meetings, and colloquia. Ei ENGINEERING MEETINGS covers all disciplines of engineering including civil engineering, environmental engineering, geological engineering, bioengineering, and electrical engineering. Non-engineering disciplines covered include electronics, control devices and principles, applied mathematics, and physics. Many application-oriented references were obtained through this on-line file search.

The second stage of the literature search was a manual one. This enabled us to follow leads generated by the online searches, and to review the many articles and books not captured by the automated systems. Much of the manual search was conducted at the National Bureau of Standards Library, supplemented by the University of Maryland Engineering and Physical Sciences Library, and the University of Maryland Computer Science Programs Library. The most productive sources were the International Journal of Man-Machine Studies, the various Proceedings of the International Joint Conference on Artificial Intelligence, the more recent Proceedings of the American Association for Artificial Intelligence, various Proceedings of the International Conference Cybernetics and Society, and the Proceedings of Trends and Applications 1983 on Automating Intelligent Behavior: Applications and Frontiers.

## 1.3 ORGANIZATION OF THE REPORT

Although the chapters in this report address interrelated issues, they have been designed to be self-contained. The reader is directed to related discussions by suitable cross-references. Where appropriate, the chapters are divided into sections, each with a discussion of a significant issue or subject matter of the chapter. The references and abstracts included in a chapter are the most relevant to the subject matter and issues of the chapter. The reader should recognize that the placement of a reference (and abstract) in a particular chapter is a judgment call of the editors, often backed up by the authors' statements; again, where necessary, items are cross-referenced.

The report also contains a listing of all sources; this should provide a useful starting place for those interested in furthering the search and obtaining complete copies of an item. We also include a glossary of artificial intelligence terms; and finally, a table of abbreviations and their meaning is given.

# 2. THE NATURE OF EMERGENCY MANAGEMENT

## 2.1 UNDERSTANDING THE NATURE OF EMERGENCIES

Knowing the nature of an emergency is the sine qua non of an effective response. However, response is only one element, albeit a major one, of the totality of actions that are generated by an emergency. Understanding the full implications of an emergency will aid in directing both passive and dynamic actions that can be applied during the mitigation, preparedness and recovery phases, as well as during the actual response.

### The Many Dimensions of an Emergency

By dimensions of an emergency we mean the various functional aspects that have to be considered during an actual emergency situation. These items are given in Figure 2.1 and range from forewarning to societal response, see FEMA [1983]. In general, they include the effects and consequences that result from the emergency. For example, given a major emergency such as the 8.1 Richter scale earthquake that hit Mexico City, we must deal with having no ample forewarning, widespread damage due to collapsed buildings and fires, loss of communications, injuries and deaths, loss of homes, and the need to search and rescue as soon as possible. During such occurrences, normal activity is at a standstill; business is not as usual. The community must pull together, be supplied with the essentials, and begin to look to the future. Other dimensions of the emergency must now be addressed: assessment of economic losses, restoration and reconstruction, aiding the victims and the community to return to normal times. Similar functional dimensions apply to other types of emergencies, as shown in Figure 2.1. The understanding of these dimensions with respect to their implications for a specific emergency forms the basis of effective emergency planning.

In our literature review, we noted a number of excellent discussions that dealt with the many aspects of commonly occurring emergencies, see Bolt [1978], White and Haas [1975], and Whittow [1979]. We also found material on some often overlooked, but important dimensions that characterize almost every emergency. We next briefly discuss a few of the more important ones.

With every emergency, there is a need of a societal response, with the character of such responses varying by emergency. A basic and important type of societal response is that of citizen volunteer groups. Although these groups tend to coalesce and operate during and after a disaster, they are also valuable entities in the mitigation and preparedness phases of an emergency. Citizen groups have organized to develop local evacuation plans, train and prepare a community on how to respond to an earthquake, and rebuild damaged areas. It is important for emergency managers to recognize the importance of these groups to the overall management of an emergency, and to work with and cooperate with them so as to be able to channel their activities in the best direction. An excellent discussion on the role of emergent citizen groups and their implications for emergency management is given in Stallings and Quarantelli [1985].

Figure 2.1 Emergency Dimensions

| EMERGENCY | Forewarning | Search & Rescue | Financial Losses | Deaths & Injuries | Evacuation | Media Management | Medical Emergency | Loss of Communication | Transportation Requirements | Information Requirements | Social Response |
|---|---|---|---|---|---|---|---|---|---|---|---|
| Nuclear plant disaster | N | | • | • | • | • | • | | • | • | • |
| Earthquake | N | • | • | • | • | • | • | • | • | • | • |
| Hurricane | • | • | • | • | • | • | • | • | • | • | • |
| Floods | • | • | • | • | • | • | • | • | • | • | • |
| Volcanoes | • | • | • | • | • | • | • | • | • | • | • |
| Enemy attack | • | | • | • | • | • | • | • | • | • | • |
| Civil disorder | • | • | • | • | | • | • | | | • | • |
| Terrorist attack | N | • | • | • | | • | • | • | | • | • |
| Epidemic | • | | | • | • | • | • | | | • | • |
| Power failure | N | | | | • | | | | | • | • |
| Resource shortage | • | | • | | | • | | | | • | • |
| Economic emergency | • | | • | | | • | | | | • | • |
| Agricultural emergency | • | | • | | | • | | | | • | • |
| Boating, ships, ferry accidents | N | • | • | • | | | • | • | • | • | |
| Airplane crash | N | • | • | • | | | • | | | • | N |
| Fire | N | • | • | • | • | • | • | | • | • | • |
| Explosion | N | • | • | • | | | • | | | • | N |
| Mine disaster | • | • | • | • | • | • | • | • | • | • | • |
| Hazardous spills | N | | | • | • | • | • | | • | • | N |
| Toxic airborne release | N | | | • | • | • | • | | • | • | N |
| Transportation disaster | N | • | • | • | | • | • | | • | • | |
| Dam failure | • | • | • | • | • | • | • | | • | • | • |
| Water shortage/pollution | • | | • | • | | • | • | | | • | • |
| Air pollution | N | | | • | • | • | • | | • | • | N |

N – Negligible

The presence of the mass media is to be expected in all emergency situations. Although the mass media sometimes appear to exacerbate and distort the events of an emergency, their cooperation and actions can be of great value in the proper management of all phases of an emergency. For example, the media can assist in pre-disaster education, serve as an effective part of the early warning system, provide information and advice to victims, and help stimulate an effective response to a disaster. However, in contrast, the demands of the media can interfere with the actions of the emergency managers, may spread rumors that contradict the reality of the disaster, and can sometimes be an unjust postdisaster critic. Given these observations, emergency planners cannot neglect the role of the media in all phases of emergency management. The "act of plugging the media into a disaster plan" involves considerable planning to achieve the required cooperation. For a discussion on how to cope with the media based on actual experience (in Canada) see Scanlon et al. [1985].

The proper financing of emergency management activities is an important consideration for all levels of government. At the local level, emergency managers must be knowledgeable about all types of financial assistance. This is especially important in those situations when little or no State and Federal aid is forthcoming. Such aid, when available, can take on many forms and the cognizant emergency managers must be able to interpret the regulations in a manner that causes the best package of benefits to be applied to their localities. Although little research has been done in this area, the reader is directed to Kunreuther and Miller [1985] and Settle [1985].

The final emergency dimension to be discussed is that of the special nature of the medical response to a disaster. There are many ways in which such medical activities differ from that of the "usual" daily emergency situations. Emergency medical personnel have few opportunities to respond to disasters and thus, are not able to routinize and practice the special needs required by extreme situations. During a disaster, medical personnel must adapt their many-on-one training and experience to be able to assess quickly the extent of an injury, and to render basic emergency care to a large number of persons under hardship conditions. Also, during a disaster, medical organizations must work with and coordinate their efforts with outside agencies. These aspects of emergency medical response emphasize that disasters are not just intensified everyday emergencies, but must be planned for in a special manner, see Tierney [1985] for further discussion on this topic.

## Natural vs. Man-Made/Technological Emergencies

In this section we examine some important differences between natural and technological emergencies and the additional dimensions that need to be considered in managing the latter type.

Man-made/technological disasters such as nuclear reactor accidents, hazardous chemical spills, and toxic airborne releases have become a major concern for emergency managers. As compared to natural disasters for which there is a substantial experience on response procedures, emergency managers are often unfamiliar with specific technological hazards. Although most local emergency units have received some technological disaster training, hands on experience is still rare. Technological disasters usually occur quickly and without warning, for example, a highway chemical spill. The short time between the onset of the damaging effects of such disasters compounds the problems of the already difficult response requirements. The resulting hazards may not be readily observable and may require the expertise of specialists for identifying, estimating and evaluating the risks.

Citizen response to a technological disaster may differ with respect to the response given a natural disaster. With a natural disaster (like a flood), communities are generally familiar with the threat and have experienced actual incidents. Also, local governments may have instituted controls to prevent or reduce the damage. Given the low incident and random nature of technological disasters, a given community, suddenly experiencing such an emergency, would probably have little knowledge about the possible resulting hazards and proper response procedures. The dependence on governmental management of an emergency is far greater in the case of a technological one. Further discussion is contained in the articles by Kasperson and Pijawka [1985] and Zimmerman [1985].

### Common Dimensions Among the Various Emergencies

The management of emergencies is a practical matter. It requires emergency managers to consider their activities in terms of cost, efficiency and effectiveness. Although each type of emergency tends to be unique, there are similarities among them, the recognition of which would help to increase the efficiency of the total emergency management structure. One way of examining such similarities is through the dimensions of emergency response given in Figure 2.1. For example, although they are required due to different causes, the response dimensions of hurricanes and earthquakes are identical (except for forewarning). There are, of course, differences in type and magnitude of response elements, especially between technological and natural disasters. It is suggested here that by understanding the similarities and differences in the response dimensions that exist for the collective set of responses faced by a community, better and more effective use of emergency-related resources can be made.

## 2.2 THE PURPOSE OF EMERGENCY MANAGEMENT

Stated very simply, the primary purpose of managing any emergency is to alleviate and/or prevent the adverse effects and consequences of a disaster. Implied by this definition is a myriad of response activities to be carried out by an organization that is structured to handle a wide range of disasters. Further, as emergency management must be viewed as an ongoing process, we must include the organization's activities during the phases of emergency mitigation, preparedness and recovery.

In a typical emergency response, lives and property have to be protected, communication links established, emergency personnel and supplies transported, search and rescue operations carried out, medical assistance provided, and so on. Mitigative measures may include hazard analysis, resource allocation, public information programs, legislating for and implementation of building and land use codes. The activities of preparedness complement those of mitigation; they include establishing a warning system, operation of a disaster monitoring center, organizing and training of emergency response teams, establishing resource management and evacuation plans. The recovery phase includes the activities of a disaster assistance center, economic assessment, cleanup and disposal, rebuilding, counseling, review and modification of codes. This totality of emergency activities demonstrates the complexity of emergency management. And, when we consider the range of emergencies that may arise, each with its special dimensions and consequences, we begin to appreciate the full nature of the emergency management task. See Drabek [1985], Perry and Nigg [1985], Seigel [1985], and Cohen and Noll [1977] for discussion of the various phase activities; and Krzyztofowicz and Davis [1984], Graff [1985] and Cohen, Noll and Weingast [1977] for discussion on the activities for managing some specific types of emergencies.

## The Scope of Emergency Response

An important aspect of emergency response is the amount of time and resources that must be spent in managing it and the range of agencies involved. This is usually a function of the size and scope of the emergency. One way of classifying the magnitude of an emergency is by ranking them according to whether their scope is local, regional, national, or even international, see Figure 2.2 and FEMA [1983]. This type of categorization only aids in anticipating the extent of involvement; any individual emergency may not fall into a strict category. As the actual size and scope of many disasters become apparent only after they have struck, it is important for local emergency management organizations to have ongoing communication channels with other local, regional and Federal response agencies.

## Mobilization Categories in Emergency Management

Mobilization is an important element of any emergency response. The type of mobilization is dependent on the type of resources required to manage an emergency, see FEMA [1983]. Mobilization planning should also take the timing of management actions into account. The following seven mobilization categories are based on the type of resources required to manage emergencies:

(1) Military mobilization is the act of preparing for war or other emergencies through assembling and organizing military resources. It is the process by which the armed forces or part of them are brought to a state of readiness for war or other national emergency.

Figure 2.2 Scope of Emergencies

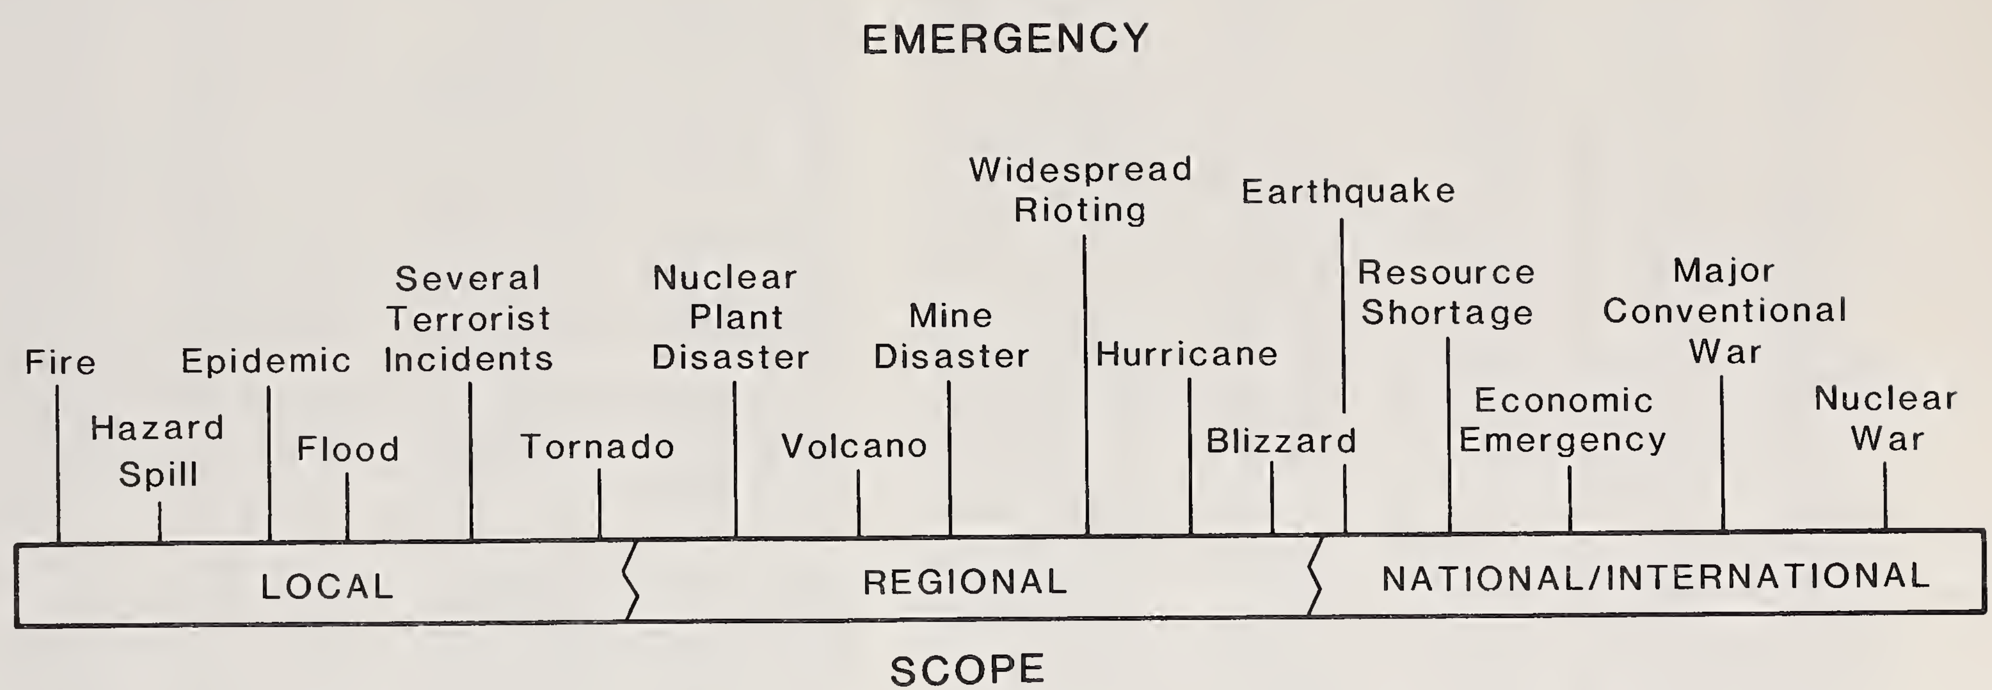

(2) Industrial mobilization is the process of marshaling the industrial sector to produce goods and services, including construction, required to support military operations and needs of the civil sector during domestic or national security emergencies.

(3) Economic mobilization is the process of marshaling the money, credit, and taxes needed to finance the management of the emergency, maintain a stable economy, and stimulate key sectors of the economy.

(4) Infrastructure mobilization is the process of marshaling the output of infrastructure systems to support the entire mobilization. Infrastructure systems include transportation, energy, communications, automated information processing, water, and agriculture.

(5) Human resources mobilization is the process of marshaling people to provide needed labor. Human resources mobilization involves identifying and allocating human resources among competing demands (see Siegel [1985] for a discussion on human resource development for emergency management).

(6) Government mobilization is the process of marshaling resources of Federal, State, and local governments to carry out the tasks required to manage emergencies. It involves bringing to the appropriate state of readiness the leadership, policymaking groups, legislative bodies, and courts; and supporting communications, facilities, procedures, and authorities to manage the emergency. Government mobilization activates and controls other aspects of mobilization.

(7) Civil preparedness mobilization is the process of marshaling resources to provide protection for the people, industry, and institutions of the United States against the effect of the spectrum of emergencies. Civil mobilization involves providing for warning and emergency instructions to the public; relocation of people to safe areas; shelter, food, water, medical care, and other human needs; and recovery and reconstruction following the emergency.

The timing of management actions for mobilization can be viewed under the framework of four phases, based on the types of activities predominantly carried out during each of these phases. These activities correspond to mitigation, preparedness, response, and recovery, respectively. These stages provide a useful framework for planning mobilization, and, in reality, these stages may not be as clear cut.

In Phase 1, Normal Operations Phase, there is a base level state of activity which exists in the absence of an overt decision to take extraordinary measures. In this phase, mitigation actions may be directed to prevent or reduce the probability of occurrence of an emergency and minimize the potential for damage or loss of life.

Phase 2, Preparedness Phase, involves taking preparedness actions to implement plans, procedures, and programs. It begins with a decision by appropriate authorities to increase the readiness of the Nation in one or more areas for a major domestic emergency or during periods of escalating international tensions.

Phase 3, Emergency Phase, involves a domestic or national security emergency, up to and including conventional or nuclear war. During this phase, response actions are taken directly before, during, and, after the onset of the emergency, to end the emergency so as to protect the public, provide critical assistance, limit damage, and reduce the probability of secondary effects.

Phase 4, Recovery Phase, involves restoring systems to normal. During this phase, short-term recovery actions are taken to assess damage and return vital life-support systems to minimum operating standards; long-term actions may continue for many years.

## 2.3 TECHNOLOGY AND NON-HUMAN RESOURCES IN EMERGENCY MANAGEMENT

### The Current Need for Technology

Each emergency requires a different understanding about the cause-effect relationships of the resulting situation, specific planning procedures, involvement of response personnel, and so on. The number of elements that have to be analyzed and responded to, even for a small-scale emergency, can become quite trying to an emergency manager. The manager must make decisions that are based on timely and accurate information; information that is gathered, filtered, and transmitted by the most suitable technological means. This may include computer-based communications that gather automatic sensor readings, data analysis and statistical programs, as well as resource and logistics management procedures.

### The Impact of Technology on Emergency Management

Technological aids to emergency decision making such as early warning systems and hurricane tracking have improved the emergency management process. Large amounts of emergency data are now gathered, analyzed and transmitted by electronic links. Computers have been used for simulations in preparedness training (Clymer [1983] and Carroll [1983]), for earthquake impact assessment (Russo and Wilson [1983] and Vance [1983]), for evacuation planning (McLean et al. [1983]), and for mobilization planning (Balbach, Fittipaldi and Webster [1983]). Various computer-based management information systems and, more recently, decision support systems have been developed as aids to emergency management (Wallace and DeBalogh [1985]). The advent of microcomputers has made the widespread use of such decision aids possible. This trend will continue in the future and increases the opportunity for technological improvements in emergency management.

An important technological possibility is the power of the computer to analyze a diverse set of information with respect to an expert body of knowledge contained within the computer's memory. This concept is embodied in the term expert system. An emergency manager -- and even an emergency management team -- cannot be expected to have the full range of expertise required to manage properly an emergency situation. Computer-based expert systems, designed to replicate the information processing and decision structures of emergency management specialists, represent a possible major improvement in the decision making abilities of emergency managers.

## References

Balbach, H. E., J. J. Fittipaldi, and R. D. Webster, "The Mobilization Facilities Planning System: A Simulation Program for Emergency Planning," Proceedings of the Conference on Computer Simulation in Emergency Planning, Vol. 11, No. 2, La Jolla, CA: The Society for Computer Simulation, Jan. 1983.

Bolt, B. A., Earthquakes: A Primer, San Francisco: W. H. Freeman and Company, 1978.

This book explores the latest theories of earthquake causes, prediction, and environmental effects. It also provides a basis for understanding the attempts to forecast the strength of earthquakes. In addition to background information on structural geology and faults, the theory of plate tectonics is presented to help explain the causes and pattern of earthquakes around the world. The relation between earthquakes and active geological faults, volcanoes, and tidal waves is discussed. The author points out the implications of the theory of plate tectonics for forecasting earthquakes, along with what actually occurs while an earthquake is taking place, and suggestions for reducing earthquake hazards.

Burby, R. J., and S. P. French, "Coping with Floods: The Land Use Management Paradox," Journal of American Planning Association, Vol. 3, July 1981.

Butler, J. E., Natural Disasters, Victoria, Australia: Heinemann Educational Australia, 1979.

This book discusses the scientific explanations of various natural disasters and their effect on the human and physical environment. Through investigating causes, it explores means of prediction, prevention, and control.

Carroll, M. J. "How to Turn an Emergency Simulation into a 'Video Game' for Training," Proceedings of the Conference on Computer Simulation in Emergency Planning, Vol. 11, No. 2, La Jolla, CA: The Society for Computer Simulation, Jan. 1983.

Clary, B. B., "The Evolution and Structure of Natural Hazard Policies," Public Administration Review, Vol. 45, Special Issue, Jan. 1985.

This article discusses the various natural hazard policies and related actions categorized according to stages of emergency management. Among the conclusions drawn by the author are: (1) a significant proportion of actions in recent years at all levels of government has been regulatory in intent, (2) many states and localities have acted on their own in developing innovative approaches to hazard control, (3) available evidence indicates that improvements can be made in emergency preparedness planning, especially where intergovernmental and/or interagency coordination is required, (4) disaster aid has become part of the cost-control question, and (5) natural hazards is a policy area in transition, and much greater emphasis is being given to pre-disaster planning and mitigation.

Cohen, L., and R. Noll, "The Economics of Disaster Defense: The Case of Building Codes to Resist Seismic Shock," Social Science Working Paper No. 130, California Institute of Technology, Pasadena, CA, Mar. 1977.

The purpose of this paper is to provide a preliminary exploration into the problem of applying economic analysis to the design of building codes in earthquake-prone areas. The purpose of seismic building codes is to increase the resistivity of buildings and other structures to major earthquakes. For the most part, seismic codes apply to new buildings and specify certain minimum design features that new buildings must match or beat. Usually these codes are designed to enable a structure to withstand ground motions associated with violent earthquakes without collapsing. In some instances, new information about the vulnerability of structures or the nature of the threat of an earthquake in a particular area leads local governments to adopt codes for existing buildings as well as new structures. For example, the Field Act in California, passed after the damaging Long Beach earthquake of 1934, required that all schools in the state, old and new, be built to withstand a major earthquake. In the aftermath of the 1971 San Fernando earthquake, measures have been adopted or proposed that would require retrofitting of particularly hazardous structures, such as dams or public meeting halls and theaters, to bring them up to the standards required of new structures.

Cohen, L., R. Noll, and B. Weingast, "Responses to Disaster: Planning for a Great Earthquake in California," Social Science Working Paper No. 131, California Institute of Technology, Pasadena, CA, Apr. 1977.

Comfort, L. K., "Integrating Organizational Action in Emergency Management: Strategies for Change," Public Administration Review, Vol. 45, Special Issue, Jan. 1985.

This article examines the role of information search processes within and between organizations as a means of integrating multiple agency and/or jurisdictional operations into effective emergency response. Creating effective organizational response under the complex, uncertain operating conditions of a major disaster poses a sobering challenge to public service agencies which bear the primary responsibility for emergency management. The perceived responsibilities of national, state, and local agencies in a disaster vary with political and economic conditions. Finding the optimal mix of shared responsibilities within the specific limitations of time, resources, and professional skills available in any given emergency, constitutes a continuing task for emergency service personnel. The author recommends that the following courses of action be considered: (1) conduct a rigorous research project on the effects of introducing a preliminary model of concurrent information search processes into the organizational design of the emergency management system, (2) revise the proposed model of concurrent information search processes in emergency management on the basis of the research findings and to implement this revised model in limited, but carefully monitored conditions for actual service delivery, (3) review the findings from the demonstration projects in terms of their applicability to the full range of interorganizational operations in emergency management, and (4) continue to refine the model in scheduled practice, with systematic implementation, observation, and evaluation of results to achieve an effective, professional, integrated system of emergency management.

Congressional Research Service, "Information Technology for Emergency Management, Report Prepared for the Subcommittee on Investigations and Oversight," Library of Congress, Washington, DC, Oct. 1984.

Drabek, T. E., "Managing the Emergency Response," Public Administration Review, Vol. 45, Special Issue, Jan. 1985.

The thesis of this article is that efforts to promote integrated emergency management information system (IEMIS), or any other disaster management tool, require explicit recognition of fundamental structural qualities that characterize emergency responses within American society. These are discussed under the sections (a) the American scene; (b) who actually responds?; and (c) strategies for enhancing intergovernmental coordination. The author offers the following six implications for managing emergency responses: (1) search and rescue actions following large-scale disasters in American society reflect the structural qualities of localism, lack of standardization, unit diversity, and fragmentation, (2) the demand structures produced by disasters exert pressures for episodic transformations of their internal operating systems and their external patterns of relationships with other agencies, (3) event qualities, type and extent of planning, and the pattern of interpersonal relationships among the organizational managers form a web of constraints that structure these emergent response networks, (4) while local governmental agencies form the core of such response networks, federal and state agencies perform crucial functions, (5) federal and state agencies have specialized roles within such response systems to which they bring unique expertise and resources in many forms, and (6) much of the strength within the American system stems from its decentralized quality, and such loosely coupled systems contrast sharply with certain business firms wherein product standardization derives from tight control structures, environmental stability, and rigid application of Taylor's principles of scientific management.

Erley, D., and W. J. Kockelman, Reducing Landslide Hazards: A Guide for Planners , Chicago: American Planning Association, 1981.

Federal Emergency Management Agency, Federal Preparedness Circular 2, Federal Emergency Management Agency, Washington, DC, 1983.

Godschalk, D. R., and D. J. Brower, "Mitigation Strategies and Integrated Emergency Management," Public Administration Review, Vol. 45, Special Issue, Jan. 1985.

This article examines present hazard mitigation practices of local agencies in terms of their potential for developing effective integrated emergency management systems. It reviews the basic concepts involved in such systems, mitigation tools available to hazard managers, and multiple hazard analysis techniques. It argues for packaging mitigation measures with development management programs into coordinated strategies in order to address effectively the opportunities and problems of integrated hazard mitigation.

Graff, G., "Beyond Bhopal: Toward a 'Fail-Safe' Chemical Industry," High Technology, Apr. 1985.

Helms, J., "Threat Perceptions in Acute Chemical Disasters: Social Aspects of Acute Chemical Emergencies," Journal of Hazardous Material, Special Issue, Vol. 4, 1981.

Hewitt, K., and I. Burton, The Hazardousness of a Place: A Regional Ecology of Damaging Events, Toronto: University of Toronto Press, 1971.

Jaffe, M., J. A. Butler, and C. Thurow, Reducing Earthquake Risks: A Planner's Guide, Chicago: American Planning Association, 1981.

Kasperson, R. E., and K. D. Pijawka, "Societal Response to Hazards and Major Hazard Events: Comparing Natural and Technological Hazards," Public Administration Review, Vol. 45, Special Issue, Jan. 1985.

This article inquires into the range of problems encountered by society as it attempts to avoid and respond to the hazardous events rooted in technology. Two tasks in particular are recognized: first, to characterize the hazard management process and to highlight the particularly difficult problems encountered, and second, to assess how the major hazardous events arising from technology affect people, their communities, and their institutions. Among the conclusions drawn by the authors are: (1) the major burden of hazard management in developed societies has shifted from risks associated with natural processes to those arising from technological development and application, (2) technological hazards pose different, and often more difficult management problems than do natural hazards because of their unfamiliarity and newness, (3) hazards may be conceived as composed of a set of linked stages begining with human needs and ending in adverse consequences, (4) in idealized form, hazard management consists of four major activities, namely, assessment, control analysis, selection of management strategy, and implementation and evaluation, (5) technological disasters tend to elicit a different pattern of public response than to natural disasters, including a greater reliance upon governmental authorities and a reduced use of community and family social networks, and (6) the emergence of a therapeutic community to ameliorate effects during the post-disaster period appears substantially less likely for technological than natural disasters.

Krzysztofowicz, R., and D. R. Davis, "Toward Improving Flood Forecast-Response Systems," Interfaces, Vol. 14, No. 3, May-June 1984.

To save lives and reduce property damage, the National Weather Service (NWS) provides flood forecasts and warnings serving 3,000 flood-prone areas. A comprehensive plan for improving this large-scale service has been formulated by the NWS Office of Hydrology based on several systems-economic studies. The one described highlights: (1) the concept of a flood forecast-response system which captures interactions among hydrologic, organizational, behavioral, and economic factors; (2) the concept of performance measures which capture the interplay between the quality of forecasts and the quality of floodplain dwellers' response, and (3) the economic value of flood forecasts and the benefits to be derived by the general public from the planned improvements.

Kunreuther, H., and L. Miller, "Insurance Versus Disaster Relief: An Analysis of Interactive Modelling for Disaster Policy Planning," Public Administration Review, Vol. 45, Special Issue, Jan. 1985.

This paper is concerned with ways that policy makers can plan ahead with respect to dealing with consequences of natural hazards. The use of decision-aiding approaches to policy planning is illustrated by focusing on trade-offs between insurance and disaster relief programs. The authors also develop a framework for incorporating a decision support system as an integral part of the policy-making process, and apply it to the case of insurance/disaster relief policy options.

Kusler, J. A., "Liability as a Dilemma for Local Managers," Public Administration Review, Vol. 45, Special Issue, Jan. 1985.

The discussion focuses on the broad range of legal issues in hazards management, which is an increasingly important issue for local governments. The legal issues and trends in the courts, including constitutional challenges to government issuance or denial of permits, planning, zoning, and mitigation measures are also reviewed. Finally, some strategies for reducing potential liability are suggested.

Martell, D. L., et al., "An Evaluation of Forest Fire Initial Attack," Interfaces, Vol. 14, No. 5, Sept.-Oct. 1984.

To help forest fire managers of the Ontario Ministry of Natural Resources evaluate initial attack resources, a relatively simple simulation model of the kinds of fires that occur in the province of Ontario and the resources used to fight them was developed. The model was useful in planning the acquisition of air tankers and the future use of air tankers, transport aircraft, and fire fighters, for initial attack purposes.

May, P. J., "FEMA's Role in Emergency Management: Recent Experience," Public Administration Review, Vol. 45, Special Issue, Jan. 1985.

McElyea, W. D., D. J. Brower and D. R. Godschalk, Before the Storm: Managing Development to Reduce Hurricane Damages, Chapel Hill, NC: Center for Urban and Regional Studies, University of North Carolina, 1982.

McLean, M. A., et al., "CLEAR: A Model for the Calculation of Evacuation Time Estimates in Emergency Planning Zones," Proceedings of the Conference on Computer Simulation in Emergency Planning, Vol. 11, No. 2, Jan. 1983.

McLoughlin, D., "A Framework for Integrated Emergency Management," Public Administration Review, Vol. 45, Special Issue, Jan. 1985.

Mushkatel, A. H., and L. F. Weschler, "Emergency Management and the Intergovernmental System," Public Administration Review, Vol. 45, Special Issue, Jan. 1985.

This article discusses the importance of the intergovernmental policy system for emergency management. The authors point out that several components of the Federal Emergency Management Agency (FEMA) organization and mission apparently are ill-suited to the challenge the agency faces within the context of the existing intergovernmental policy system. Among the issues discussed are the major characteristics of the political system within which FEMA developed and operates, the intergovernmental setting and some fiscal considerations affecting emergency management. A policy matrix is also described which conceptually represents the challenges to FEMA in implementing an integrated emergency management system.

National Governors' Association, Comprehensive Emergency Management: A Governor's Guide. Washington, D.C.: Center for Policy Research, 1979.

Pasqualetti and D. Pijawka, (Eds.), Nuclear Power: Assessing and Managing Hazardous Technology, Boulder, CO: Westview Press, 1984.

Perry, R. W., Comprehensive Emergency Management, Greenwich, CT: JAI Press, 1984.

Perry, R. W., and A. H. Mushkatel, Disaster Preparedness: Warning, Response and Community Relocation, Westport, CT: Greenwood Press, 1984.

Perry, R. W., and J. M. Nigg, "Emergency Management Strategies for Communicating Hazard Information," Public Adminstration Review, Vol. 45, Special Issue, Jan. 1985.

The purpose of this paper is to examine the process of communicating information to the public about environmental risks to increase the likelihood that citizens will adopt protective measures. To accomplish this objective, the authors examine three issues in public education. First, an attempt is made to understand how emergency management agencies can become identified as credible sources of information within the community. Second, attention is given to establishing and maintaining viable communication channels through which local emergency management personnel can reach the public. The closing section of the paper discusses strategies for increasing citizen receptivity to officially designated prospective measures.

Petak, W. J., and A. A. Atkinsson, Natural Hazard Risk Assessment and Public Policy, New York: Springer-Verlag, Inc., 1982.

Petak, W. J., "Emergency Management: A Challenge for Public Administration," Public Administration Review, Vol. 45, Special Issue, Jan. 1985.

The author of this article argues that the concept of an integrated emergency management system is based on the belief that the efforts of many disciplines are necessary if we are to reduce the consequences of natural and man-made disasters. This implies that the public administrator, as emergency manager, must have the conceptual skill to understand: (1) the total system, (2) the uses to which the products of the efforts of various professionals will be put, (3) the potential linkages between the activities of various professional specialists, and (4) the specifications for output formats and language which are compatible with the needs and understanding of others within the total system. As with any other problem-focused management system, the emergency management system must give full consideration to the problems associated with authority relationships, conflict management (i.e., who is responsible, who is accountable, who pays, and who controls), and legal authority versus perceived responsibility.

Quarantelli, L., Delivery of Emergency Medical Services in Disasters: Assumptions and Realities, New York: Irvington Publishers, 1983.

Rubin, C. B., Managing the Recovery From a Natural Disaster, Management Information Service Report, Volume 14, Washington, DC: International City Management Association, 1982.

Rubin, C. B., and D. G. Barbee, "Disaster Recovery and Hazard Mitigation: Bridging the Intergovernmental Gap," Public Administration Review, Vol. 45, Special Issue, Jan. 1985.

Russo, J. E., and R. R. Wilson, "Earthquake Vulnerability Analysis for Economic Impact Assessment," Proceedings of the Conference on Computer Simulation in Emergency Planning, Vol. 11, No. 2, La Jolla, CA: The Society for Computer Simulation, Jan. 1983.

Scanlon, J., et al., "Coping with Media in Disasters: Some Predictable Problems," Public Administration Review, Vol. 45, Special Issue, Jan. 1985.

The authors of this article propose that the media perform some very important functions in times of disasters, such as, providing pre-disaster education, serving as an effective warning system, providing information and advice to victims and others in the wake of disasters, activating the local disaster response, and stimulating effective disaster relief. The authors also develop an effective plan for media relations in disaster.

Schmidt, W. J., et al., "Tri-purpose Emergency Management Model," Proceedings of the Conference on Emergency Planning, Volume 15, No. 1, La Jolla, CA: The Society for Computer Simulation, Jan. 1985.

The authors of this article propose that an effective emergency manager must consider three operational objectives: (1) his plan must provide a logical, flexible progression of response actions, (2) personnel involved in laboratory execution must understand the logic and be trained in its execution, and (3) the plan must be tested under a variety of scenarios to evaluate its effectiveness. They describe an initial version and an improved version of a system model that was developed by the Management

Systems Laboratories of Virginia Tech. The initial version was applied in plan development and analysis for the Department of Energy's Savannah River Plant.

Selman, V., A. L. Selman, and J. Selman, "Earthquakery/Quackery: The Science/Art of Predicting Natural Disasters," Proceedings of the Conference on Emergency Planning, Volume 15, No. 1, La Jolla, CA: The Society for Computer Simulation, Jan. 1985.

Settle, A. K., "Financing Disaster Mitigation, Preparedness, Response and Recovery," Public Administration Review, Vol. 45, Special Issue, Jan. 1985.

This paper examines the sources of aid and financial devices that are available to local governments during disasters. There is a lack of information on financing emergency management and particularly on funding mechanisms that can prepare local government to pay for disasters when no federal or state aid is forthcoming. Because of limited tax base and reserves, community governments' options of financing disasters is very restricted. The author provides a matrix on funding alternatives connected to the different phases of emergency management. The matrix is particularly helpful in determining what sources and amounts of revenue would be available to a community to pay the costs of an emergency when no federal or state emergency is declared.

Siegel, G. B., "Human Resource Development for Emergency Management," Public Administration Review, Vol. 45, Special Issue, Jan. 1985.

Stallings, R. A., and E. L. Quarantelli, "Emergent Citizen Groups and Emergency Management," Public Administration Review, Vol. 45, Special Issue, Jan. 1985.

In this paper the authors address three principal questions: (a) What are the characteristics of emergent citizen groups in disasters? (b) What roles do they play in the various phases of the emergency management process? and (c) What implications do such groups have for emergency management policies and procedures? Drawing primarily from research by the Disaster Research Center (DRC) on emergent social groupings in actual or potential disasters, supplemented by the authors' own interpretation of selected empirical studies of others, they suggest the following implications for emergency management: (1) the formation of emergent citizen groups is inevitable before, during, and after disasters, (2) emergent citizen groups are the outcome of natural social processes; they do not represent a deviational or abnormal pattern, (3) emergent citizen groups, while not always functional, are not necessarily dysfunctional either, (4) emergent citizen groups are not inherently in opposition to public authorities, and (5) emergent phenomena cannot be eliminated by prior planning. They also suggest how these characteristics of citizen emergent groups can be harnessed during emergency management operations.

Tierney, K. J., "Emergency Medical Preparedness and Response in Disasters: The Need for Interorganizational Coordination," Public Administration Review, Vol. 45, Special Issue, Jan. 1985.

The author of this paper observes that unlike more common medical emergencies, large-scale disasters are infrequent; therefore, emergency responders do not have an opportunity to practice and repeat disaster-related tasks until they become routine. The infrequent nature of disasters also means that the need for preparedness is obscured; pressing daily problems get priority. Disasters are also distinctive because they require health care providers to perform in ways that are contrary to their training and experience. In providing medical care, it is suggested that the emergency managers should consider the following points in attempting to develop more successful disaster emergency management system (EMS) preparedness programs: (1) the problem of delivering EMS to victims of disasters and other mass casuality events in the United States is not a problem of insufficient resources and technology or an inadequate supply of skilled service providers, (2) persons with responsibility for emergency management at different levels of government need to exchange information with local, regional and state EMS officials and providers, as well as with public and private health care organizations, (3) medical sector groups need to work more closely with the major disaster preparedness and response organizations to develop realistic and compatible disaster response procedures, and (4) EMS disaster preparedness planning should be an element in comprehensive local, regional, state, and federal preparedness programs, not a separate effort.

U.S. Department of Commerce, Inventory of Natural Hazards Data Resources in the Federal Government, Boulder, CO: National Geophysical and Solar-Terrestrial Data Center, May 1979.

Vance, E. D., "Earthquake Hazard Assessment Utilizing Geographic Information System Integration," Proceedings of the Conference on Computer Simulation in Emergency Planning, Vol. 11, No. 2, La Jolla, CA: The Society for Computer Simulation, Jan. 1983.

Walker, H., "Spatial Data Requirements for Emergency Response," Proceedings of the Conference on Emergency Planning, Volume 15, No. 1, La Jolla, CA: The Society for Computer Simulation, Jan. 1985.

Wallace, W. A., and F. de Balogh, "Decision Support Systems for Disaster Management," Public Administration Review, Vol. 45, Special Issue, Jan. 1985.

In this article the authors focus on an increasingly important component of information technology, decision support systems (DSS), and describe its role in disaster management. Specifically, they (1) differentiate DSS from more traditional management information systems (MIS); (2) discuss the disaster life cycle and its stages in the context of a typology of decisions -- strategic, tactical, and operational -- found within those stages; and (3) illustrate the potential of DSS technology with some selected applications in the area of earthquake mitigation, crisis response, toxic chemical spill control, and emergency staff preparedness training. They also review some future technological developments in this area.

White, G. F., and E. J. Haas, Assessment of Research on Natural Hazards, Cambridge, MA: The MIT Press, 1975.

Whittow, J., Disasters: The Anatomy of Environmental Hazards, Athens, GA: The University of Georgia Press, 1979.

Zimmerman, R., "The Relationship of Emergency Management to Government Policies on Man-Made Technological Disasters," Public Administration Review, Vol. 45, Special Issue, Jan. 1985.

The two main questions addressed in this paper are: (a) To what extent do emergencies or disasters exist due to man-made technologies involving hazardous materials? (b) What mechanisms are currently in place to cope with emergencies, and how well are they adapted to them? To address the former, the author uses trend data on the origins or releases of contaminants, health effects, and transport of contaminants to show whether a gradual or sudden accumulation of adverse conditions exists. The second question is treated in the framework of one application of an emergency management typology for natural hazards to existing legislation and programs for the control of non-natural technological hazards so as to gauge their orientation toward emergencies.

# 3. INTRODUCTION TO EXPERT SYSTEMS

An expert system can be defined as a computer-based model that is able to replicate the decisions of one or more human experts in a well-defined subject area. Furthermore, the types of problems addressed by expert systems usually require significant human expertise for their solution. An expert system models or mimics the expertise by which the human expert solves the problem. The major components that serve to model a human decision maker's expertise are referred to as the knowledge base and the inference procedure.

The knowledge base records rules, facts, and information about the problem that may be useful in formulating a solution. The inference procedure provides a control structure for analyzing the information contained in the knowledge base (i.e., facts and rules).

It is the distinct, separate nature of the knowledge base and the inference procedure which together captures the human decision maker's expertise. This separation and its associated flexibility is absent in traditional decision-aiding models which use a rigidly defined control structure to manipulate the information contained in a database. Knowledge is the key; the performance of an expert system is, to a large degree, a function of the size and quality of its knowledge base. The term knowledge-based expert system is used to emphasize the importance of knowledge in modeling human expertise.

## 3.1 SELECTED EXPERT SYSTEMS CURRENTLY IN USE

Over the past decade, expert systems have been developed for a number of different problem domains. Several of the better-known applications are listed in Table 3.1. In all, information on seven problem domains is listed; an example of an expert system for each domain is also given. Table 3.1 is intended to provide an indication of the variety of such systems currently in use. Additional information on expert system applications is given in Chapter 8 and the papers by Gevarter [1982] and Nau [1983].

From Table 3.1, we can see that the problem domain tends to be focused on well-defined subjects (e.g., chemical analysis and computer system design). With the exception of R1/XCON, used by the Digital Electronics Corporation for customizing its VAX computer systems, and DENDRAL, which is used extensively by researchers in chemistry, most systems have remained in the laboratory stage of development. Hopefully, this research orientation will change due in part to the availability of better tools for building expert systems. These tools should significantly reduce the high cost of system development. Furthermore, it is likely that many of the second generation of expert systems will be developed specifically for commercial applications, rather than to explore topics which are primarily of a research nature.

Table 3.1 Some Existing Expert Systems Classified by Problem Domain

| Name | Problem Domain | Selected System Purpose |
|---|---|---|
| DENDRAL | Interpretation | To generate plausible structural representations of organic molecules from mass spectogram data. |
| MYCIN | Diagnosis | To diagnose bacterial infections and make recommendations for antibiotic therapy. |
| R1/XCON | Design | To configure VAX computer systems based on a customer's order of components. |
| VM | Monitoring | To monitor and interpret physiological data of patient breathing. |
| Drilling Advisor | Troubleshooting | To determine the likely causes of drilling problems and recommend curative and preventive treatments. |
| STEAMER | Instruction | To instruct users how to operate a steam propulsion plant. |
| Spill Crisis | Control | To interpret, recommend and monitor crisis management activities. |

## 3.2 THE STRUCTURE OF AN EXPERT SYSTEM: AN IDEALIZATION

We next present an overview of an "ideal" expert system. This overview is abstracted from the text Building Expert Systems. The ideal expert system is shown schematically in Figure 3.1. We first begin with a capsule summary of the ideal expert system. The various components of the system are then described in greater detail and related to specific entries in Figure 3.1. An ideal expert system contains: (1) a language processor for problem-oriented communications between the user and the expert system, (2) a global data base for recording intermediate results, (3) a knowledge base comprising facts, as well as planning and problem-solving rules, (4) a control structure or inference procedure, which includes an interpreter that applies the rules, a scheduler to control the order of rule processing, a consistency enforcer that adjusts previous conclusions when new data (or knowledge) alter their bases of support, and (5) a justifier that rationalizes and explains the system's behavior. The text which follows defines the terms introduced above; interrelationships are shown through reference to Figure 3.1.

### The Language Processor

The user interacts with the expert system via a problem-oriented language. The language processor mediates information exchange between the user and the expert system. Typically, the language processor interprets the user's questions, commands, and volunteered information. Conversely, the language processor formats and sends to the user information generated by the system, answers to questions, explanations and justifications for its behavior, and requests for data. Several key topics governing the user interface are given in Hayes-Roth [1984].

### The Knowledge Base

The knowledge base of an expert system consists of facts and rules. The facts constitute a body of information that would be generally agreed upon by experts in the field. The rules are often based upon less well-known heuristics (rules of plausible reasoning, rules of thumb, rules of good guessing) that characterize expert-level decision making in the subject area. Further information on what constitutes a knowledge base and how it is constructed in practice is given in: Buchanan and Duda [1983], Davis [1982], and Michaelson, Michie and Boulanger [1985].

### The Global Data Base

The global data base records the user's inputs, intermediate hypotheses, and decisions. Every expert system has some form of intermediate decision representation. However, only a few explicitly employ all of the components shown in Figure 3.1. In addition to the user's inputs, three types of elements are identified as being included in the global data base: (1) plan; (2) agenda; and (3) solution elements. Plan elements describe the overall or general attack the system will pursue against the current problem, including current plans, goals, problem states, and contexts. The agenda elements

Figure 3.1 Schematic of an Ideal Expert System

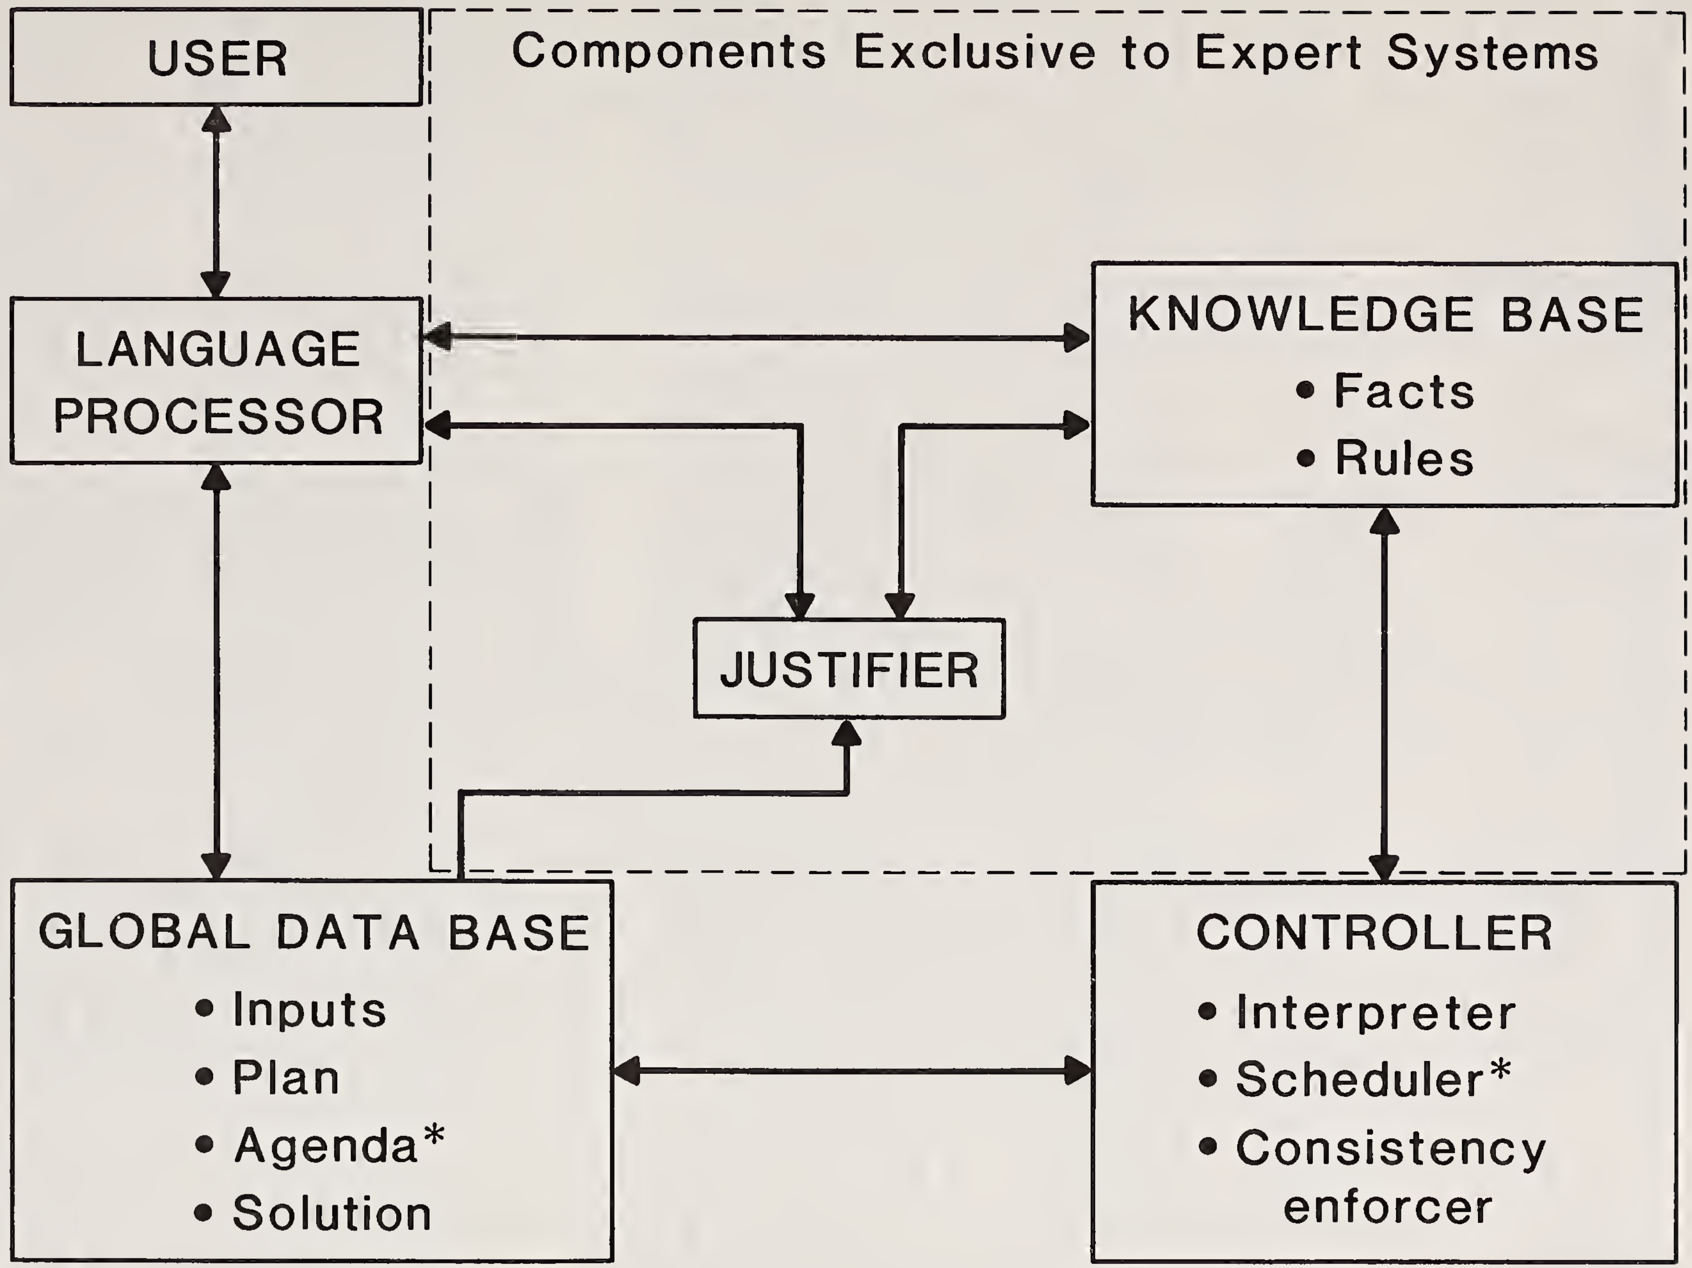


*These items are unique to expert systems record the potential actions awaiting execution, which generally correspond to knowledge base rules that seem relevant to some decision previously placed in the global data base. The solution elements represent the candidate hypotheses and decisions the system has generated thus far, along with the dependencies that relate decisions to one another. For a more general interpretation of the global data base, see Gevarter [1982] and Nau [1983].

## The Controller

The control structure determines how to use the rules contained in the knowledge base; it contains three elements: (1) an interpreter, (2) a scheduler, and (3) a consistency enforcer. The interpreter executes the chosen agenda item by applying the corresponding knowledge base rule. Generally, the interpreter ensures the relevance conditions of the rule, binds variables in these conditions to particular global data base solution elements, and then makes those changes to the global data base that the rule prescribes. The scheduler maintains control of the agenda and determines which pending action should be executed next. Schedulers may embody considerable knowledge, such as doing the most profitable thing next. To apply such knowledge, the scheduler needs to give each agenda item a priority according to its relationship to the plan and other extant solution elements. To do this, the scheduler generally needs to estimate the effects of applying the potential rule. The consistency enforcer attempts to maintain a consistent representation of the emerging solution. Most expert systems use some kind of numerical adjustment scheme to determine the degree of belief in each potential decision. Such schemes attempt to ensure that plausible conclusions are reached and inconsistent ones are avoided. For an indepth discussion of control structures, see Nilsson [1980] and Buchanan and Duda [1983].

## The Justifier

Finally, the justifier explains the actions of the system to the user. In general, it answers questions about why some conclusion was reached or why some alternative was rejected. To do this, the justifier uses a few general types of question-answer plans. These typically require the justifier to trace backward along solution elements in the global data base from the questioned conclusion to the intermediate hypotheses or data that support it. Each step backward corresponds to the inference of one knowledge base rule. The justifier collects these intermediate inferences and translates them into English for presentation to the user.

There are a number of important distinctions between expert systems and conventional computer-based models which are highlighted by the dashed rectangle in Figure 3.1. The key distinctions are the inclusion of the knowledge base and a justifier. Additional comparisons between expert systems and conventional computer-based models are given in Nau [1983] and Quinlan [1980].

## 3.3 TOOLS AND TECHNIQUES FOR BUILDING AND TESTING AN EXPERT SYSTEM

One of the primary reasons for the increased use of expert systems is due to the tools which are available to the analyst. To better understand the capabilities provided by these tools, it is useful to discuss their role in the building of an expert system at two levels. The first level focuses on the traditional artificial intelligence programming languages. The second level focuses on what have become known as shells. A shell may be thought of as a specialty package which facilitates the construction of the knowledge base and provides the inferencing procedure for exploiting the information contained within the knowledge base. Both types of tools are discussed in greater detail in Chapter 5.

If one were to build an expert system, one might choose an artificial intelligence programming language such as LISP or PROLOG. Within the United States, LISP (List Processor) is the artificial intelligence programming language of preference. Since LISP is a symbol manipulation language, it promotes the representation of knowledge which many authors claim to be oriented more towards symbolic processing than numeric processing.

PROLOG (Programming in Logic) is widely used in Europe and Japan. The motivation behind PROLOG is to permit the analyst to specify the tasks in terms of logic statements, rather than specify how they should be processed numerically by the computer. One difficulty with the use of artificial intelligence programming languages is the current lack of standardization. Consequently, many researchers have been moving toward the use of more universally accepted languages such as FORTRAN and PASCAL, see Thompson and Thompson [1985].

Since many applications share common methods for representing knowledge, performing inference, and maintaining large systems, researchers have designed shells which facilitate the development of expert systems. The motivation for developing a shell is to provide support to the analyst in representing symbolic knowledge. This knowledge can include facts, definitions, heuristic judgments, and procedures for doing a task or achieving a goal. The shell may be programmed in an artifical intelligence language so as to exploit its symbolic manipulation capabilities or a common language. The main benefit in building an expert system by using a shell is that the analyst is relieved of the painstaking task of programming in languages such as LISP or FORTRAN. Consequently, using the shell enables the analyst to spend more time formulating rules or developing heuristics and less on the production of code. The use of shells is therefore an important step in reducing the costs of developing an expert system.

Despite the promise of expert systems, the current state-of-the-art limits their capabilities. The expert systems have rather narrow domains of expertise, partially because there are no efficient means for building and maintaining large knowledge bases. Extraction and representation of knowledge is performed by knowledge engineers. The knowledge engineer must cope with the expert's inability to impart knowledge in a manner that can be easily encoded, while the problem domain expert needs to understand the system's problem solving methods and the representational framework. In almost all expert systems, the explanation of its line of reasoning is nothing more than the system retracing its steps to the rules that enabled the system to draw that conclusion. The expert systems often provide little guidance to the user about their appropriateness for new problems or about their scope and limitations. As expert systems are not able to deal with incompatible knowledge items, expertise is usually limited to that of a single individual. However, it is often desirable to merge expertise from several experts. Expert systems are not efficient in handling many types of reasoning concepts such as causal reasoning, reasoning over time, scene (or 3-dimensional) representation, and analogical reasoning. For example, reasoning over time requires dynamic representation and a mechanism for handling feedback.

The above discussion points out some of the areas in which expert systems must undergo improvements and characteristics future expert systems must incorporate. The advent of parallel processing is likely to increase the scope of expert systems by easing the constraints imposed by the complexity and size of the knowledge base and the tasks performed. There is also a move toward building what are called the deep-knowledge systems that reason by using knowledge of structure, function, and behavior of objects, see Kinnucan [1984]. These systems are expected to function better than the current expert systems especially in those domains where novel problems are frequent. As a remedy for the current (predominantly rule-based) knowledge representation technique's inadequacy in handling concepts such as causal reasoning, scene representation, etc., there is an increasing trend towards building systems using semantic networks and frames. Expert systems will also benefit as systems for understanding speech and visual images are perfected. These advances will open up many problem domains in which expert system technology can be applied.

## REFERENCES


Barnett, J. A., and M. I. Bernstein, "Knowledge-Based Systems: A Tutorial," Report No.: SDC-TM-(L)-5903/000/00, System Development Corporation, Santa Monica, CA, 1977.

Barr, A., and E. A. Feigenbaum, (Eds.), The Handbook of Artificial Intelligence (Vols. I and II), Los Altos, CA: Kaufman, 1981-1982.

Brachman, R. J., et al., "What are Expert Systems?" in F. Hayes-Roth, D. A. Waterman, and D. B. Lenat (Eds.), Building Expert Systems, Reading, MA: Addison-Wesley Publishing Company, 1983.

In this article the authors tackle the issue of definitions and boundaries to explain what an expert system is and why. They point out that several traits characterize an expert system, including the symbolic nature of the task it performs, a broad and robust intelligence, an ability to rationalize and justify its behavior, a capacity to expand its range of capabilities and refine its skills, and an ability to solve important problems involving complexity and uncertainty.

Bramer, M. A., "A Survey and Critical Review of Expert Systems Research," in D. Michie (Ed.), Introductory Readings in Expert Systems, New York: Gordon Breach Science Publishers, 1982.

Bramer, M. A. (Ed.), Research and Development in Expert Systems: Proceedings of the Fourth Technical Conference of the British Computer Society Specialist Group on Expert Systems, Cambridge, England: Cambridge University Press, 1984.

Buchanan, B. G., "New Research on Expert Systems," in J. E. Hayes, D. Michie, and Y. H. Pao (Eds.), Machine Intelligence 10, Chichester, England: Horwood, 1982.

Buchanan, B. G., and E. H. Shortliffe (Eds.), Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project, Reading, MA: Addison-Wesley Publishing Company, 1984.

Buchanan, B. G., and R. O. Duda, "Principles of Rule-Based Expert Systems," in M. Yovits (Eds.), Advances in Computers, Volume 22, New York: Academic Press, 1983.

The authors describe an expert system as a computer program that provides expert-level solutions to important problems and is: (1) heuristic, i.e., it reasons with judgmental knowledge, as well as with formal knowledge of established theories, (2) transparent, i.e., it provides explanations of its lines of reasoning and answers to queries about its knowledge, (3) flexible, i.e., it integrates new knowledge incrementally into its existing store of knowledge. The purpose of this article is to familiarize the readers with the architecture and construction of the important class of rule-based systems. Many of the existing expert system programs, or issues involved are not dealt with, but an attempt is made to provide a framework for understanding this advancing frontier of computer science.

Chester, D., "Elements of Knowledge-Based Expert Systems," Proceedings of Micro-Delcon '82, The Delaware Bay Computer Conference, March 1982.

Expert systems are built to solve problems in application areas for which "good" algorithms are not known. These systems consist of a global data base of assertions, a set of rules that represent small bits of an expert's knowledge, and a control strategy for applying the rules to the assertions. Agendas are used to make the control strategy more efficient. Efficiency can be further increased by indexing the rules and assertions in various ways, one of which is frames. A system for deriving formal specifications from natural language requirements is presented as an example.

Cohen, P., and E. A. Feigenbaum, (Eds.), The Handbook of Artificial Intelligence (Vol. III), Los Altos, CA: Kaufmann, 1982.

D'Ambrosio, B., "Expert Systems - Myth or Reality?" Byte, Vol. 9, No. 1, Jan. 1985.

David, H., "An Analysis of Expert Thinking," International Journal of Man-Machine Studies, Vol. 18, 1983.

Davis, D. B., "Assessing the Strategic Computing Initiative," High Technology, Apr. 1985.

Davis, R., "Expert Systems: Where are We? and Where Do We Go From Here?" AI Magazine, Vol. 3, No. 2, 1982.

Duda, R. O., and E. H. Shortliffe, "Expert Systems Research," Science, Vol. 220, No. 4594, Apr. 1983.

Research in artificial intelligence (AI) has several goals. One is the development of computational models of intelligent behavior, including both its cognitive and perceptual aspects. A more engineering-oriented goal is the development of computer programs that can solve problems normally thought to require human intelligence. This article concerns a class of AI computer programs intended to serve as consultants for decision-making. They are often called "expert systems" because they address problems normally requiring human specialists for their solution. Some of these programs have reached expert levels of performance on the problems for which they were designed. The authors describe these accomplishments as well as identify some difficult problems that must be solved in order to realize their benefits in practice.

Duda, R. O., and J. G. Gaschning, "Knowledge-Based Expert Systems Come of Age," Byte, Vol. 6, No. 9, Sept. 1981.

Dutta, A., and A. Basu, "An Artificial Intelligence Approach to Model Management in Decision Support Systems," Computer, Vol. 17, No. 9, Sept. 1984.

Feigenbaum, E., and P. McCorduck, "The Fifth Generation, An American Response," Computerworld, Vol. 17, No. 20, May 1983.

The Japanese are at present developing a revolutionary generation of computers to be known as knowledge information-processing systems, or kips. The real power of such systems lies in their capacity to reason with enormous amounts of information. The Japanese expect them to be in general worldwide use by the 1990s. The authors examine the heart of this fifth generation, the applied side of artificial intelligence called knowledge-based systems. They trace their history and describe the applications of expert systems in fields such as medicine, chemistry, and defense. The authors stress that it is unacceptable for the USA to have to depend on Japan for vital defense technology, and summarize the options open to Americans. They examine in detail one proposal for a National Centre for Knowledge Technology.

Gaines, B. R., and M. L. G. Shaw, "Expert Systems and Simulation," in G. Birtwistle (Ed.), Artificial Intelligence, Graphics, and Simulation, La Jolla, CA: Simulations Councils, Inc., 1985.

Simulation and expert systems (ES) are contrasted as computer-based technologies for the support of human knowledge processes. An overview is given of ES developments emphasizing that an ES is not just a static encoding of expertise but rather a dynamic process promoting the acquisition of knowledge. Applications of ES's to industrial control problems are discussed where the links to classical control and simulation can be most readily seen. Simulation and ES's are shown to be complementary technologies. The combination of simulation and ES technologies is examined and shown to form a basis for very powerful system designs.

Gevarter, W. B., Intelligent Machines: An Introductory Perspective of Artificial Intelligence and Robotics, Englewood Cliffs, NJ: Prentice-Hall, Inc., 1985.

Gevarter, W. B., "An Overview of Expert Systems," Report No.: NBSIR 82-2505, National Bureau of Standards, Gaithersburg, MD, May 1982.

This report provides an overview of the rapidly evolving field of expert systems -- what an expert system is, the techniques employed, existing systems, applications, who is doing it, who is funding it, the state-of-the-art, research requirements, and future trends and opportunities.

Hall, H. K., J. C. Moore, and A. B. Whinston, "A Theory of Expert Systems," Proceedings of the First Symposium on the Theory and Application of Expert Systems in Emergency Management Operations, Special Publication 717, Gaithersburg, MD: National Bureau of Standards, 1986.

Harmon, P., and D. King, Expert Systems, New York: Wiley Press, 1985.

Hayes-Roth, F., "Knowledge-Based Expert Systems: A Tutorial," Computer, Vol, 17, No. 9, Sept. 1984.

This article presents an overview of the field of knowledge engineering. It describes the major developments that have led to the current great interest in expert systems, then presents a brief discussion of the principal scientific and engineering issues in the field, as well as of the process of building expert systems, the role of tools in that work, how expert systems perform human-computer interface functions, and the frontiers of research and development.

Hayes-Roth, F., "Knowledge-Based Expert Systems," Computer, Vol. 17, No. 10, Oct. 1984.

Problem-solving engines organize the activity of knowledge systems to solve problems. To understand these engines, we need to relate their implementation to their design and intended purpose. Today's knowledge systems aim to solve specific problems. A knowledge engineer analyzes the problem to be solved and then adopts an overall approach generally consisting of (1) a problem-solution paradigm, (2) a general knowledge-system architecture, and (3) a specific problem-solving. Today's problem-solving engines provide specific devices for implementing the knowledge engineer's choices.

Hayes-Roth, F., D. A. Waterman, and D. B. Lenat, "An Overview of Expert Systems," in F. Hayes-Roth, D. A. Waterman, and D. B. Lenat (Eds.), Building Expert Systems, Reading, MA: Addison-Wesley Publishing Company, 1983.

In this paper the authors trace the principal developments in artificial intelligence that have led to the current emphasis on knowledge-based expert systems and the corresponding field of knowledge engineering. Machines that lack knowledge seem doomed to perform intellectually trivial tasks. Those that embody knowledge and apply it skillfully seem capable of equaling or surpassing the best performance of human experts. Knowledge provides the power to do work; knowledge engineering is the technology that promises to make knowledge a valuable industrial commodity.

Keller, R., and P. Townsend, "Knowledge-Based Expert Systems," Computerworld, Vol. 18, No. 49A, Dec. 1984.

An overview of knowledge-based, or expert systems is given. The authors discuss application generators and the features of knowledge-based products. An outline of research activities in the business area is given.

Kinnucan, P., "Knowledge-Based Systems are Tackling Problems that Once Required Human Expertise," High Technology, Jan. 1984.

Kurzweil, R., "What is Artificial Intelligence Anyway?" American Scientist, Vol. 73, May-June 1985.

The controversy surrounding AI is evidenced by a certain amount of discord within the field. Academic researchers have accused some industrial AI developments of being shallow and giving too high a priority to short-term commercial goals rather than long-term research. Industrial researchers have accused academic AI of providing superficial demonstrations, of not developing robust systems that really work. Companies have accused their competitors of not using real AI techniques and so on. Unfortunately we are not likely to find published any time soon the definitive work that will lay this confusion to rest. The author approaches these questions and the central question "What is artificial intelligence anyway?" by discussing what the discipline of artificial intelligence is now capable of achieving, what it will be capable of achieving, and how we can best reach those goals.

Lee, R. M., "On Information System Semantics: Expert vs. Decision Support Systems," Society of Science and Information Studies, Vol. 5, No. 1, Jan. 1985.

AI will surely have an important impact on the management of future organizations. The problem now is to make some reasonable assessment of its promise and limitations in order to direct research and development. Here we examine what an ideal knowledge-based management information system could and could not do. The arguments are based on considerations of formal semantics.

Lee, R. M., "Expert vs. Management Support Systems: Semantic Issues," Cybernetics and Systems, Vol. 14, No. 2-4, Apr.-Dec. 1983.

Expert systems hold great promise for technical application areas such as medical diagnosis or engineering design. They are, however, less promising for management applications. The reason is that managers are not experts in the sense of possessing a formal body of knowledge which they apply. The limitations of artificial intelligence approaches in managerial domains is explained in terms of semantic change, with attention toward management (decision) support systems.

Martins, G. R., "The Overselling of Expert Systems," Datamation, Nov. 1984.

McCartney, J., "Expert Systems: Expertise Captured in a System," Data Processing, Vol. 24, No. 3, Mar. 1982.

Michaelsen, R. H., D. Michie, and A. Boulanger, "The Technology of Expert Systems," Byte, Vol. 10, No. 4, Apr. 1985.

This is an introductory level article on expert systems. It also discusses some methods for building them, including the advantages and disadvantages of each method. Finally, the article reviews the computer resources needed to build and run expert systems.

Michie, D., "Expert Systems," The Computer Journal, Vol. 23, No. 4, 1980.

Miller, R. K., Artificial Intelligence: A New Tool for Industry and Business, Volume 1 (Technology and Applications), Madison, GA: SEAI Institute, 1984.

Moore, T., "Artificial Intelligence: Human Expertise from Machines," EPRI Journal, June 1985.

Nau, D. S., "Expert Computer Systems," Computer, Vol. 16, No. 2, Feb. 1983.

The author points out the main difference between most expert systems and ordinary computer applications programs: in the former, the model of problem-solving in the application domain is explicity in view as a separate entity or knowledge base rather than appearing only implicitly as part of the coding of the program. This knowledge base is manipulated by a separate, clearly identifiable control strategy. Such a system architecture provides a convenient way to construct sophisticated problem-solving tools for many different domains. Further, ordinary computer programs organize knowledge on two levels: data and program. Most expert computer systems, however, organize knowledge on three levels: data, knowledge base, and control. Computers organized in this way are often called knowledge-based systems. This article discusses the technique used in expert systems on each of these levels. Information on the areas of AI problem-solving and knowledge-representation techniques is also included.

Naylor, C., Build Your Own Expert System, New York: Wiley Press, 1985.

Nilsson, N. J., Principles of Artificial Intelligence, Palo Alto, CA: Tioga Publishing Company, 1980.

The author attempts to fill a gap between artificial intelligence theory and practice. Unlike many books on artificial intelligence, which divide the subject into its major areas of application, namely, natural language processing, automatic programming, robotics, machine vision, automatic theorem proving, intelligent data retrieval systems, etc., this book is organized on the basis of general computational concepts involving the kinds of data structures used, the types of operations performed on these data structures, and the properties of control strategies used by AI systems. The important roles played by generalized production systems and the predicate calculus are stressed. Among the major topics covered are rule-based systems, robot problem-solving systems, and structured-object representations.

Quinlan, J. R., "An Introduction to Knowledge-Based Expert Systems," Australian Computer Journal, Vol. 12, No. 2, May 1980.

Discusses the concepts underlying knowledge-based systems, and shows how they differ from more conventional (algorithm-based) systems. Three principal dimensions are identified: The representation of knowledge, architectures for deploying knowledge, and techniques for acquiring knowledge. Key ideas are illustrated from a selection of current expert systems covering a wide range of applications.

Rich, E., Artificial Intelligence, New York: McGraw-Hill Book Company, 1983.

Sauers, R., and R. Walsh, "On the Requirements of Future Expert Systems," Proceedings of the Eighth IJCAI, American Association for Artificial Intelligence, 1983.

Many artificial intelligence applications require the use of expert systems. As expert systems move into new domains, several significant changes can be expected. Among these are an increase in number of rules in the rule base and in the number of data elements contained in working memory. Also, many new applications require that expert systems move into real-time domains. Here, a system must be able to process large quantities of data which are changing rapidly. Many problem-solving situations will be time critical, and the system must take into account the availability and distribution of scarce system resources. For these reasons, the efficiency of a given expert system design and its ability to perform complex memory management tasks will become increasingly important. This will require modifications in the traditional production system architectures. In this paper, the design requirements of future expert systems are discussed, and HAPS, a recently implemented production system architecture designed to address these issues, is presented.

Schindler, M., "Expert Systems," Electronic Design, Vol. 33, No. 1, Jan. 1985.

This technology forecast discusses the research being undertaken into artificially intelligent systems. Because qualified personnel are in extremely short supply, many researchers in artificial intelligence laboratories are already trying to come up with tools and methodologies to simplify the development of future AI products. Languages, knowledge base schemas, user interfaces, and inference paradigms all are undergoing reevaluation with an eye on efficiency. Another obstacle must be overcome before a reliable AI industry can develop: The hardware on which its programs run must become cheaper. That is why artificial intelligence companies are addressing such mundane problems as execution speed, code size, and transportability, things that AI workers have always loved to ignore.

Sen, A., "Decision Support System: An Expert Systems Approach," Proceedings of the 16th Annual Meeting of the American Institute for Decision Sciences, Atlanta, GA: American Institute for Decision Sciences, Nov. 1984.

Focuses on the design aspects of a decision support system (DSS). Although DSS have traditionally been designed to address a specific problem domain, the author believes current expert system design techniques can be exploited to develop domain-independent, user-friendly systems. The development of such DSS, called the XDSS, is described.

Shoben, A. M., "Expert Systems: Newest Brainchild of Computer Science," Report No.: RAND/P-6552, RAND Corp., Santa Monica, CA, June 1981.

Document discusses an emerging field of computer science called expert systems. These are computer programs that seek to combine expert knowledge about a domain (in theory, any domain) with expert methods of conceptualizing and reasoning about that domain. The decision-making power of such programs rests upon a knowledge base that puts together factual information about the domain with the heuristics (informal rules of thumb) experts use to rapidly find solutions to problems. This specialty in turn is part of artificial intelligence -- the science, as one definition has it, of making machines do things that people need intelligence of a high order to do. Despite its small size and relative newness, the expert systems field is attracting widespread interest because of its demonstrated usefulness in medicine and chemistry, and for its potential application to such diverse and complex problem areas as structural engineering, military strategy, air traffic control, crisis management and many others.

Shurkin, J. N. "Expert Systems: The Practical Face of Artificial Intelligence," Technology Review, Nov.-Dec. 1983.

Sloman, A., "Epistemology and Artificial Intelligence," in D. Michie (Ed.), Expert Systems in the Micro-Electronic Age, Edinburgh: Edinburgh University Press, 1979.

A favorite area of disagreement among AI researchers concerns the selection of a formalism for specifying computational mechanisms. Many of the systems are based on the idea of a set of production-rules, namely rules with a condition and an action. The underlying mechanism is, in the simplest cases, a machine which repeatedly looks for a rule whose condition matches some data-base, and then obeys the action. This paper attempts to supplement the detailed studies of expertise in the other conference papers with some general remarks about epistemological issues arising out of artificial intelligence. The author points out that these issues are not unique to the design of clever artifacts, but are central to the study of the human mind.

Stefik, M., et al., "The Organization of Expert Systems: A Tutorial," Artificial Intelligence, Vol. 18, 1982.

Tate, P., "The Blossoming of European AI," Datamation, Nov. 1984.

Thompson, B. A., and W. A. Thompson, "Inside an Expert System," Byte, Vol. 10, No. 4, Apr. 1985.

In this article the authors describe one way that an expert system can use a set of rules to conduct a consultation session. A description of the mechanics of how an expert system operates by creating a "cardboard inference engine" is given, and some of the programming considerations for translating the cardboard system into PASCAL are examined.

Vaucher, J. G., "Views of Modeling: Comparing the Simulation and AI Approaches," in G. Birtwistle (Ed.), Artifical Intelligence, Graphics, and Simulation, La Jolla, CA: Simulations Councils, Inc., 1985.

Researchers and practitioners in the field of simulation and those in artificial intelligence (AI) have had to face quite similar problems in creating models of complex and sometimes partially understood systems. To a large extent, solutions have been developed independently in each area, leading to techniques and software tools which differ markedly in terminology but often overlap in terms of concepts. The recent stress on knowledge representation in AI has emphasized a common ground, modeling of reality, but each group maintains a slightly different emphasis: dynamic behavior for simulationists and logical inference for AI workers. In this paper, modeling tools and practice in both areas are contrasted and useful areas of cross-fertilization are suggested.

Waldrop, M. M., "The Necessity of Knowledge," Science, Vol. 223, March 1984.

Waltz, D. L., "Artificial Intelligence," Scientific American, Vol. 247, No. 4, Oct. 1982.

Computer programs written by investigators in artificial intelligence have demonstrated conclusively that in certain activities (including activities most people would say require intelligence, such as playing games) the computer can outperform a human being. Recent programs have demonstrated that the computer can even develop elaborate theories about a limited domain such as arithmetic from a few simple axioms. At the same time the understanding of various features of human intelligence has been considerably enriched by the attempt to describe analogues of those features in the detail necessary for writing a program. As a result, the analogy relating the performance of the computer to that of human intelligence has broadened and matured. This article is a good introduction to the kinds of results that can be achieved by programs in artificial intelligence. Most of the programs that are discussed are well-known approaches to major problem areas in the field. An attempt is also made to clarify the main contributions of each of these programs.

Waterman, D. A., A Guide to Expert Systems, Reading, MA: Addison-Wesley Publishing Company, 1985.

Winston, P. H., Artificial Intelligence, Reading, MA: Addison-Wesley Publishing Company, 1984.

Yaghmai, N. S., and J. A. Maxin, "Expert Systems: A Tutorial, "Journal of American Society for Information Science, Vol. 34, No. 5, Sept. 1984.

Expert systems are intelligent computer applications that use data, a knowledge base, and a control mechanism to solve problems of sufficient difficulty that significant human expertise is necessary for their solution. Expert systems use artificial intelligence problem-solving and knowledge-representation techniques to combine human expert knowledge about a problem area with human expert methods of conceptualizing and reasoning about that problem area. As a result, it is expected that such systems can reach a level of performance comparable to that of a human expert in a specialized problem area. The high-level knowledge base and associated control mechanism of expert systems are in essence a model of the expertise of the best practitioners of the problem area in question and, hence, human users are provided with expert opinions about problems in that area. Expert systems do not pretend to give final or ultimate

conclusions to displace human decision making; they are intended for consulting purposes only.

Yazdani, M., and A. Narayanan (Eds.), Artificial Intelligence: Human Effects, Chichester, England: Ellis Horwood Limited, 1984.

The main theme of this book is to explore the interaction between the fast-growing discipline of artificial intelligence (AI) and other human endeavors. Among the main topics discussed are the usefulness of AI in medicine, law, and education. Other topics include intelligent machines and human society, AI methodology, and philosophical implications of AI. Taken together, these sections examine the implications that intelligent computing systems have on our legal system, medical care, and education, as well as the effects of sharing our environment with such systems. The conclusion introduces some uncertainties concerning AI, and a postscript presents an entertaining fictional view of the future with intelligent machines.

# 4. REASONING AND METHODOLOGICAL ISSUES IN THE DESIGN AND DEVELOPMENT OF EXPERT SYSTEMS

In the previous chapter, the major components of an ideal expert system were defined. A brief description of the relationship among these components was used to illustrate how an expert system operated. The purpose of this chapter is to describe in greater detail several methodological issues which are the basis of an expert system. The focus of this chapter is on the three components of an expert system and their associated elements that deal with: (1) knowledge representation; (2) the control structure; and (3) the manipulation of data.

## 4.1 KNOWLEDGE REPRESENTATION

Knowledge representation in expert systems usually relies on one of four techniques: (1) production rules; (2) first order logic; (3) frames; or (4) semantic networks.

Production rules are by far the most common means for knowledge representation in expert systems. An expert system that uses production rules is termed a rule-based expert system. A production rule consists of two parts, an antecedent and a consequent. The antecedent represents some pattern of information and the consequent specifies an action to be taken when the data matches the pattern. The antecedent typically contains several clauses linked by the logical AND and OR connectors. The consequent consists of one or more verb phrases that specifies the action to be taken. In the IF-THEN terminology frequently used to describe production rule sytems, the antecedent is the IF clause and the consequent is the THEN clause.

First order logic is a formal way of representing logical propositions and relations between propositions. First order logic is highly useful in knowledge representation because one can derive any fact that follows logically from the propositions it represents.

Frames are another technique used in knowledge representation. A frame is a data structure through which all knowledge about a particular object or event is stored. Frames include attribute data; they can also contain rules to express more complex facts or relationships, to specify default values for pieces of information about an object when that information is not explicitly given, and to invoke or link to other frames when applicable data conditions or patterns are met. In this method, the organization and structure of the knowledge is used to constrain and guide the program to a smaller and more tightly bounded search space of possible candidate solutions.

Semantic networks are another way to represent knowledge. They are similar to frames in the sense that knowledge is organized around objects and events. In this case, the objects are represented by nodes in a graph and the relations among them are represented by labeled arcs.

Each of the techniques just described have certain advantages which must be considered during the design and development of an expert system. Production rules are especially useful for representing procedural knowledge (e.g., methods for accomplishing goals). First order logic supplies a means for explicitly expressing virtually any type of knowledge. Frame-based sytems can concisely store an immense amount of knowledge about object properties and relations. Semantic networks are good for representing relations between objects.

## 4.2 CONTROL STRUCTURES

Expert systems employ various control structures or strategies to generate solutions. The control strategy is often referred to as the "inference engine." The two-types of control strategies employed are state-space search and problem reduction. We first describe state-space search procedures, followed by discussions of the key techniques of backward chaining and forward chaining. Finally, the pros and cons of problem reduction procedures are given.

In most expert system applications, the precise series of steps necessary to solve a problem are generally not known in advance. It is therefore necessary to search through a space containing many alternate paths, not all of which lead to solutions. Expert systems employ a variety of techniques to limit and constrain the number of candidate solutions that are examined. Work in the areas of artificial intelligence and operations research have produced several procedural/algorithmic approaches (e.g., depth-first, breadth-first, branch and bound) that can be applied to these search problems. It is important to point out, however, that these procedural/algorithmic-based approaches result in a certain amount of back-tracking, which has an adverse effect on the speed with which an expert system can establish a solution. This is because procedural/algorithmic approaches explore one or more paths as far as possible before returning to the search at some other point within the state space. Because of the extremely large search spaces characteristic of many expert system applications, it is necessary to further constrain and prune the search. This is accomplished through the use of heuristics that are based on informal knowledge (e.g., production rules) about the problem to be solved. These programs or operators are pattern invoked in the sense that they are not called by other programs in the ordinary way, but instead are activated by the control structure whenever certain conditions in the data are met.

Backward chaining is the most commonly used control strategy; it explores a series of hypotheses in the order of their prior likelihood. Each hypothesis is explored in turn, until all hypotheses have been exhausted. The expert system then ranks the hypotheses by the degree to which they are supported by the evidence. Backward chaining is also referred to as a goal-driven search, since the hypotheses described above may be thought of as a hierarchy of goals, subgoals and data.

Forward chaining focuses on the initial evidence contained in the global data base. The system then considers each item of evidence in turn, attempting to work forward to a goal. Forward chaining is also referred to as a data-driven or antecedent-driven search. It is important to note that some types of forward chaining will not terminate if the initial evidence does not lead to a conclusion. For this reason, many expert systems use backward chaining as a control strategy.

Many expert systems augment the control strategies just described by propagating constraints. In this approach, the set of candidate solutions become further and further constrained by rules or operators that produce local constraints on what small pieces of the solution must look like. More and more rule applications are made until no more rules are applicable and only a few candidate solutions are left. This process can be thought of as a type of state-space search that avoids backtracking, since every candidate solution must satisfy all of the constraints produced by the rule applications.

One alternative to state-space search is an approach known as problem reduction. Here, the problem to be solved is partitioned or decomposed into subproblems that can be solved separately, so that the combined solutions to the subproblems will yield a solution to the original problem. The approach may be illustrated graphically through reference to a problem-reduction graph (i.e., an AND/OR graph). All possible decompositions of the problem are shown on the graph. Within the graph, each OR branch represents a choice of several alternative decompositions, and each AND branch represents a particular way of decomposing a problem. Since some decompositions of a problem may lead to solvable subproblems and others may not, care must be exercised to ensure that a "feasible" decomposition results. The graph of such a decomposition is referred to as a solution graph.

## 4.3 THE MANIPULATION OF DATA

If one thinks of data manipulation within an expert system in the same way as in a conventional program, one would expect the tools for building expert systems to provide a variety of data structures (e.g., records, sets and arrays). As we shall see in Chapter 5, many tools for building expert systems provide a single powerful representation. This approach permits the user to focus on the problem without the burden of attending to a variety of syntax and other regulations which are inherent to conventional data-base management systems. One should also make a distinction between a knowledge base and a data base. Whereas data bases have a predetermined structure, knowledge bases consist of a set of unstructured, almost isolated facts and rules. The paths by which facts and rules are related in a knowledge base are pattern invoked (i.e., determined as needed to solve a particular problem). The relationships between data items in relational or network data bases, by contrast, are designed into the data base. Another critical distinction between the two is that data bases store all of their information explicitly, whereas the bulk of the information in a knowledge base is inferred from a few basic facts using inference rules.

## References

Aikins, J. S., "Representation of Control Knowledge in Expert Systems," Proceedings of the National Conference on Artificial Intelligence, AAAI-80, American Association for Artificial Intelligence, 1980.

Barnett, J. A., "Computational Methods for a Mathematical Theory of Evidence," Proceedings of the Seventh IJCAI, American Association for Artificial Intelligence, 1981.

Many knowledge-based expert systems employ numerical schemes to represent evidence, rate competing hypotheses, and guide search through the domain's problem space. This paper has two objectives: first, to introduce one such scheme, developed by Dempster and Shafer, to a wider audience; second, to present results that can reduce the computation-time complexity from exponential to linear, allowing this scheme to be implemented in many more systems. In order to enjoy this reduction, some assumptions about the structure of the type of evidence represented and combined must be made. The assumption made here is that each piece of the evidence either confirms or denies a single proposition rather than disjunction. For any domain in which the assumption is justified, the savings are available.

Barstow, D. R., "The Roles of Knowledge and Deduction in Algorithm Design," in J. E. Hayes, D. Michie, and Y. H. Pao (Eds.), Machine Intelligence 10, Chichester, England: Horwood, 1982.

In the earliest attempts to apply artificial intelligence techniques to program synthesis, deduction (that is, the use of a general purpose mechanism such as a theorem prover) played a central role. Recent attempts have relied almost exclusively on knowledge about programming in particular domains, with no significant role for deduction. Even in such knowledge-based systems, however, there seems to be an important role for deduction in testing the applicability conditions of specific programming rules. This auxiliary role for deduction is especially important in algorithm design, as can be seen in the hypothetical synthesis of a breadth-first enumeration algorithm. The interplay between knowledge and deduction also shows how one can be motivated to consider the central mathematical properties upon which particular algorithms are based, as illustrated in the synthesis of two minimum cost spanning tree algorithms.

Ben-Bassat, M., "Toward a Theory of Expert System: A Problem-Driven Approach," Proceedings of the COMPCON 83 Fall, Silver Spring, MD: Computer Society Press, Oct. 1983.

Existing approaches to generalize expert systems, examples EMYCIN and ARBY, are directed from the technology space to the application space. That is, given a certain technique (e.g., a rule-based engine) additional applications are identified which might also be solved by the technique. This approach forces a given application into a strait-jacket tailored by tools that were developed for another application. We propose an approach that generalizes from the problem space to the technology space, that is, the selection of a set of techniques to match the unique requirements of a given class of problems. The abstract problem space is a collection of abstract problems (e.g., evidence interpretation) which are part of many real-life applications. This space offers a typology for classifying the components of real-life applications. By the problem-to-technology approach, a given real-life application is first decomposed into its smallest problem solving tasks, which are then matched with the abstract problems that fit them best. This, in turn, provides a list of techniques that are applicable with their corresponding assumptions and limitations. Our approach is illustrated for a set of abstract problems related to situation assessment tasks such as battlefield reading, medical diagnosis, corporate assessment, potential crisis assessment, and weather forecasting.

Ben-Bassat, M., and A. Freedy, "Knowledge Requirements and Management in Expert Decision Support Systems for (Military) Situation Assessment," Report No.: 46/81; ARI-TR-576, Israel Inst. of Business, Tel-Aviv, Aug. 1983.

Situation assessment tasks (e.g., medical diagnosis) battlefield reading, corporation assessment for merger or acquisition purposes, are formulated as a general family of problem solving tasks. We characterize the generic nature of this family as a multiprespective multimembership hierarchical pattern recognition problem, identify the types of decision problems involved in the situation assessment process, and propose a unified approach for the development of situation assessment decision support systems (DSS). The focus is on knowledge representation and elicitation, although issues related to inference mechanisms, system structure and expert-machine-user interface are also discussed. The presentation is accompanied by examples from military situation assessment. However, comparable examples from medical and business applications are also cited. Many of the ideas presented here have already been implemented in the MEDAS system; and medical DSS for emergency and critical care medicine.

Bennett, J. S., "On the Structure of the Acquisition Process for Rule-Based Systems," Machine Intelligence (Infotech State-of-the-Art Report Series 9, No. 3), Maidenhead, England: Pergamon Infotech Ltd., 1981.

Blanning, R. W., "Knowledge Acquisition and System Validation in Expert Systems for Management," Human Systems Management, Vol. 4, No. 4, Autumn 1984.

The growing number of successes in constructing expert systems for such recognized professionals and specialists as doctors diagnosing and treating infectious diseases and mathematicians performing symbolic integration has given rise to suggestions that expert systems be developed for managers in order to capture their specialized knowledge in performing decision tasks in manufacturing, marketing, finance, and natural language. Concepts at any level of precision or vagueness can be handled in a consistent way. Fuzzy logic allows appropriate deductions from rules whose conditions are only partially satisfied, and facilitates the overlap between rules essential to effective performance of a knowledge-based system. Several preliminary fuzzy systems for management applications are described, and prospects for future operational systems are discussed.

Borgida, A., "Knowledge Representation as the Basis for Requirements Specifications," Computer, Vol. 18, No. 4, Apr. 1985.

For purpose of discussion, we define the requirements phase of the software life-cycle as the stage that precedes the design of software system architecture. This phase includes analysis of customer needs, as well as specification of both the functional behavior of the proposed system and the nonfunctional requirements that must be met. It is clear that, to carry out these activities, the analyst must gain an understanding of the environment in which the software will be used and the use to which the software will be put. We contend that this understanding should be expressed and recorded as a model of the environment, and that the various requirements should be expressed in relation to this model. Once we have made a case of this contention, we present some features that we believe are desirable in any language intended to support requirements engineering. In the remainder of the article, we illustrate some features of the language RLL and point out how it arose from the confluence of two streams of thought: the representation of knowledge in artificial intelligence and the traditional concern for abstraction found in software engineering. The article concentrates on language design issues.

Brent, E., "Consultant: An Expert System Using a Relational Database," AAMSI Congress 83 Proceedings, Bethesda, MD: American Association of Medical Systems and Informatics, May 1983.

CONSULTANT is a knowledge-based expert program designed to provide a supportive intelligent learning environment for designing research samples. The paper describes CONSULTANT's approach to two traditional problems in AI programming: knowledge representation and explanation. CONSULTANT uses a relational data base to store and retrieve rules.

Brown, J. S., R. R. Burton, and J. de Cleer, "Knowledge Engineering and Pedagogical Techniques in SOPHIE I, II, and III, " in D. Sleeman and J. S. Brown (Eds.), Intelligent Tutoring Systems, London: Academic Press, 1982.

Buchanan, B. G., "New Research on Expert Systems," in J. E. Hayes, D. Michie, and Y. H. Pao (Eds.), Machine Intelligence 10, Chichester, England: Horwood, 1982.

Bundy, A., et al., "Solving Mechanics Problems Using Meta-level Inference," in D. Michie (Ed.), Expert Systems in the Micro-Electronic Age, Edinburgh: Edinburgh University Press, 1979.

In this paper the authors describe a program (MECHO), written in PROLOG, which solves a wide range of mechanics problems from statements in both predicate calculus and English. MECHO uses the technique of meta-level inference to control search in natural language understanding, common sense inference, model formation and algebraic manipulation.

The authors argue that this is a powerful technique for controlling search while retaining the modularity of a declarative knowledge representation.

Burton, R. R., "Diagnosing Bugs in a Simple Procedural Skill," in D. Sleeman and J. S. Brown (Eds.), Intelligent Tutoring Systems, London: Academic Press, 1982.

Carbonell, J. G., "Experiential Learning in Analogical Problem Solving," Proceedings of the National Conference on Artificial Intelligence, AAAI-82, American Association for Artificial Intelligence, 1982.

A computational model of skill acquisition is analyzed based on extensions to an analogical problem solving method and previous AI work on concept acquisition. The present investigation focuses on exploiting and extending the analogical reasoning model to generate useful exemplary solutions to related problems from which more general plans can be induced and refined. Starting with a general analogical inference engine, problem solving experience is, in essence, compiled incrementally into effective procedures that solve various classes of problems in a more reliable and direct manner.

Cendrowska, J., and M. A. Bramer, "A Rational Reconstruction of the MYCIN Consultation System," International Journal of Man-Machine Studies, Vol. 20, No. 3, Mar. 1984.

This article presents a detailed analysis and partial reconstruction in POP-2 of the control structure of MYCIN, arguably the best-known and most significant expert system currently in existence. The aim is to aid the development of theory in the field and to assist those who wish to build their own working systems. Attention is focused, inter alia, on the production rules, the goal-directed backward chaining of rules that comprises the control structure and the parameters and context types employed, together with the data structures created during a consultation, and the system's use of certainty factors to handle uncertain information. A detailed account is given of how a typical consultation proceeds and some variants that can occur are considered. Developments to MYCIN since its original implementation in 1976 and its generalized version, EMYCIN, are also briefly described. The article identifies a number of gaps in the original reporting of MYCIN and presents a critique of a number of its features and an appraisal of the value of a MYCIN-like approach as a starting point for further expert systems development.

Chandrasekaran, B., "Expert Systems: Matching Techniques to Tasks," in W. Reitman (Ed.), Artificial Intelligence Applications for Business, Norwood, NJ: Ablex Publishing Corporation, 1984.

In this chapter an attempt is made to relate the architectures and representations for expert systems to the types of tasks for which they are appropriate. The author starts with an analysis of the features that characterize an expert system, and discuss the need for symbolic knowledge structures that support qualitative reasoning in the design of expert systems. The author considers rules, logical formulas and frames for representation of expert knowledge. In particular he provides an analysis of the multiplicity of roles that rules have played in different rule-based systems, and emphasize the need to distinguish among these roles. The author provides a theory of types of expert problem solving and argues that such a taxonomy enables one to characterize expert system capabilities and help match problems with techniques. Throughout the chapter it is emphasized that the important issue is the nature of the information processing task in a given task domain, and issues of formalisms for representation are subordinate to that basic issue.

Chandrasekaran B., and S. Mittal, "Conceptual Respresentation of Medical Knowledge for Diagnosis by Computer: MDX and Related Systems," in M. Yovits (Ed.), Advances in Computers, Vol. 22, New York: Academic Press, 1983.

The purpose of this article is not to provide another survey of the field but, rather, to describe and discuss an approach to the design of medical decision-making systems based on the notion of conceptual structures for knowledge representation. A collection of related systems that have been under development in our laboratory exemplifies this approach, but the ideas are more general than the particular systems to be described. The central system in this group of systems is called MDX, which is a diagnostic system (i.e., it attempts to classify a given case as an element of a disease taxonomy). This system interacts with two other systems during its problem solving, PATREC and RADEX, the former a knowledge-based patient data-base system that answers MDX's queries about patient data, and the latter a radiological consultant which helps MDX in the interpretation of various kinds of imaging data. Both PATREC and RADEX are invoked by MDX as needed, but MDX is in control of the overall diagnostic process.

Chandrasekaran, B., and S. Mittal, "Deep Versus Compiled Knowledge Approaches to Diagnostic Problem-Solving," Proceedings of the National Conference on Artificial Intelligence, AAAI-82, American Association for Artifical Intelligence, 1982.

Chiou, Sr., W. C., "Systems Architecture of a Remote Sensing Expert System," Proceedings of the 1983 IEEE Computer Society Workshop on Computer Architecture for Pattern Analysis and Image Database Management, Silver Spring, MD: IEEE Computer Society Press, Oct. 1983.

This study presents a preliminary systems architecture of a knowledge-based expert system which provides a powerful approach for resolving massive remote-sensing data analysis problems. It incorporates a natural language processor for user interfacing and maintains a standardized global data base. The expert system consists of three major subsystems: a global spectral database, which acts as a distributed databank for systems analysis; knowledge-based expert subsystems, which contain general and specific geological models; and spectral image processing and understanding, which provides a set of spectral image analysis methods. An executive consultant system integrates the overall simulation and synthesis for final systems utilization.

Clark, K. L., "An Introduction to Logic Programming," in D. Michie (Ed.), Introductory Readings in Expert Systems, New York: Gordon Breach Science Publishers, 1982.

Coombs, M., and J. Alty, "Expert Systems: An Alternative Paradigm," International Journal of Man-Machine Studies, Vol. 20, 1984.

There has recently been a significant effort by the AI community to interest industry in the potential of expert systems. However, this has resulted in far fewer substantial applications projects than might be expected. This article argues that this is because human experts are rarely required to perform the role that computer-based experts are programmed to adopt. Instead of being called in to answer well-defined problems, they are more often asked to assist other experts to extend and refine their understanding of a problem area at the junction of their two domains of knowledge. This more properly involves educational rather than problem-solving skills. An alternative approach to expert system design is proposed based upon guided discovery learning. The user is provided with a supportive environment for a particular class of problems, the system predominantly acting as an advisor rather than directing the interaction. The environment includes a database of domain knowledge, a set of procedures for its application to a concrete problem, and an intelligent machine-based advisor to judge the user's effectiveness and advise on strategy. The procedures focus upon the use of user generated explanations both to promote the application of domain knowledge and to expose understanding difficulties. Simple database PROLOG is being used as the subject material for the prototype system which is known as MINDPAD.

Davis, R., "Amplifying Expertise with Expert Systems," in P. H. Winston and K. A. Prendergast (Eds.), AI Business: The Commercial Uses of Artificial Intelligence, Cambridge, MA: MIT Press, 1984.

In surveying the state-of-the-art of expert systems, two important calibration points are provided by the magnitude of investment necessary to build a robust system and the stage of development reached to date by most expert systems. Data from existing efforts, though meager, seem to suggest that even in the best of cases, at least five man-years of effort are necessary before an expert system begins to perform reliably. It is also revealing to note that most expert systems to date have been developed only through the stage of construction of the basic knowledge base. Relatively few so far have progressed to the stage of extended testing, further development, and documentation.

Davis, R., B. G. Buchanan, and E. H. Shortliffe, "Production Rules as a Representation of a Knowledge-Based Consultation Program," Artificial Intelligence, Vol. 8, 1977.

Davis, R., M. Weikert, and M. Shirley, "Diagnosis Based on Description of Structure and Function," Proceedings of the National Conference on Artificial Intelligence, AAAI-82, American Association for Artificial Intelligence, 1982.

Defude, B., "Knowledge-Based Systems Versus Thesaurus: An Architecture Problem About Expert System Design," Proceedings of the Joint BCS and ACM Symposium on Research and Development in Information Retrieval, Cambridge, England: Cambridge University, 1984.

Duda, R. O., J. G. Gaschnig, and P. E. Hart., "Model Design in the PROSPECTOR Consultant System for Mineral Exploration," in D. Michie (Ed.), Expert Systems in the Micro-Electronic Age, Edinburgh, Scotland, Edinburgh University Press, 1979.

PROSPECTOR is a computer consultant system intended to aid geologists in evaluating the favorability of an exploration site or region for occurrences of ore deposits of particular types. Knowledge about a particular type of ore deposit is encoded in a computational model representing observable geological features and the relative significance thereof. The authors describe the form of models in PROSPECTOR, focusing on inference networks of geological assertions and the Bayesian propagation formalism used to represent the judgmental reasoning process of the economic geologist who serves as model designer. Following the initial design of a model, simple performance evaluation techniques are used to assess the extent to which the performance of the model reflects faithfully the intent of the model designer. These results identify specific portions of the model that might benefit from fine tuning, and establish priorities for such revisions. This description of the PROSPECTOR system and the model design process serves to illustrate the process of transferring human expertise about a subjective domain into a mechanical realization.

Duda, R. O., and R. Reboh, "AI and Decision Making: The PROSPECTOR Experience," in W. Reitman (Ed.), Artificial Intelligence Applications for Business, Norwood, NJ: Ablex Publishing Corporation, 1984.

PROSPECTOR is an expert system that was designed for decision-making problems in mineral exploration. It uses a structure called an inference network to represent the required judgmental knowledge. Once encoded, this knowledge can be used in a variety of ways. This chapter describes the use of inference networks as a language for representing and using expert knowledge, and presents examples showing its use in three different kinds of decision-making problem.

Duda, R. O., et al., "Semantic Network Representations in Rule-Based Inference Systems," in D. A. Waterman and F. Hayes-Roth (Eds.), Pattern-Directed Inference Systems, New York: Academic Press, 1978.

Englemore, R. S., and A. Terry, "Structure and Function of the CRYSALIS System," Proceedings of the Sixth IJCAI, American Association for Artificial Intelligence, 1979.

Feigenbaum, E. A., "Themes and Case Studies in Knowledge Engineering," in D. Michie (Ed.)., Expert Systems in the Micro-Electronic Age, Edinburgh, Scotland: Edinburgh University Press, 1979.

This paper examines emerging themes of knowledge engineering, illustrates them with case studies drawn from the work of the Stanford Heuristic Programming Project, and discusses general issues of knowledge engineering art and practice.

Feigenbaum, E. A., B. G. Buchanan, and J. Lederberg, "On Generality and Problem Solving: A Case Study Using the DENDRAL Program," in B. Meltzer and D. Michie (Eds.), Machine Intelligence 6, New York: American Elsevier. 1971.

Findler, V., et al., "A Module to Estimate Numerical Values of Hidden Variables for Expert Systems," International Journal of Man-Machine Studies, Vol. 18, No. 4, Apr. 1983.

In the area of strategic decision-making, the objective often is to achieve one's own goals and to prevent the achievement of the adversaries' goal. To do so, the decision-maker needs to know, as precisely as possible, the values of the relevant variables at various times. Some of these variables, the open variables, are readily measurable at any time. Others, the hidden variables, can be measured only at certain times, either intermittently or periodically. The authors have implemented a module that can act as a decision-support tool for a variety of expert systems in need of estimates of hidden variable values at any desired time. The estimation is based on generalized production rules expressing stochastic, causal relations between open and hidden variables. The quality of the estimates improves through a multi-level learning process as both the number and the quality of the rules increase. The modularity of these causal relations make incremental expansion and conflict resolution natural and easy. Restricting the set and the domain of pattern formation rules to a reasonable size makes the system effective and efficient. Finally, the system can be easily employed for distributed database applications.

Fox, M. S., "Reasoning with Incomplete Knowledge in a Resource-Limited Environment: Integrating Reasoning and Knowledge Acquisition," Report No.: CMU-RI-TR-81-3, Robotics Institute, Carnegie-Mellon University, Pittsburgh, PA, Mar. 1981.

Fox, M. S., B. Allen, and G. Strohm, "Job-Shop Scheduling: An Investigation in Constraint-Directed Reasoning," Proceedings of the National Conference on Artificial Intelligence, AAAI-82, American Association for Artificial Intelligence, 1982.

Friedland, P., "Acquisition of Procedural Knowledge from Domain Experts," Proceedings of the Seventh IJCAI, American Association for Artificial Intelligence, 1981.

Ganascia, J. G., "Using Expert System in Merging Qualitative and Quantitative Data Analysis," International Journal of Man-Machine Studies, Vol. 20, 1984.

Gaschnig, J. G., et al., "Evaluation of Expert Systems: Issues and Case Studies," in F. Hayes-Roth, D. A. Waterman, and D. B. Lenant (Eds.), Building Expert Systems, Reading, MA: Addison-Wesley Publishing Company, 1983.

Georgeff, M., and U. Bonollo, "Procedural Expert Systems," Proceedings of the Eighth IJCAI, American Association for Artificial Intelligence, 1983.

A scheme for explicitly representing and using expert knowledge of a procedural kind is described. The scheme allows the explicit representation of both declarative and procedural knowledge within a unified framework, yet retains all the desirable properties of expert systems such as modularity, explanatory capability, and extendability. It thus bridges the gap between the procedural and declarative languages, and allows formal algorithmic knowledge to be uniformly integrated with heuristic declarative knowledge. A version of the scheme has been fully implemented and applied to the domain of automobile engine fault diagnosis.

Genesereth, M. R., "Diagnosis Using Hierarchial Design Models," Proceedings of the National Conference on Artificial Intelligence, AAAI-82, American Association for Artificial Intelligence, 1982.

Goldstein, I. P., and B. Roberts, "Using Frames in Scheduling," in P. H. Winston and D. Brown (Eds.), Artificial Intelligence: An MIT Perspective (Volume I), Cambridge, MA: MIT Press, 1979.

Green, C., and S. J. Westfold, "Knowledge-Based Programming Self-Applied," in J. E. Hayes, D. Michie, and Y. H. Pao (Eds.), Machine Intelligence 10, Chichester, England: Horwood, 1982.

Hart, A., "The Role of Induction in Knowledge Elicitation," Expert Systems, Vol. 2, No. 1, Jan. 1985.

Knowledge elicitation from experts is a major problem in building expert systems. This paper summarizes some of the difficulties inherent in the process and suggests that in certain situations induction can help.

Hartley, R. T., "How Expert Should an Expert System Be?" Proceedings of the Seventh IJCAI, American Association for Artificial Intelligence, 1981.

A computer system which aids computer engineers in fault diagnosis is described. The system, called CRIB (Computer Retrieval Incidence Bank) is shown to fit into the class of pattern-directed inference systems. Emphasis is placed on the before and after phases of system generation and it is shown why, to be called an expert system, these phases are important. The forms of knowledge used in CRIB are shown to be adequate for diagnosis and yet possess little of the structural or functional knowledge of more advanced expert systems. Summaries are given of the three phases of implementation: elicitation, implementation of knowledge structures, validation and improvement. The idea of an expert system as a "model of competence" is mentioned and the transference of the system architecture to software diagnosis, using the same model, is described. There are short discussions of system performance and the nature of expert systems.

Jackson, P., and P. Lefrere, "On the Application of Rule-Based Techniques to the Design of Advice-Giving Systems," International Journal of Man-Machine Studies, Vol. 20, 1984.

Jarke, M., and Y. Vassiliou, "Coupling Expert Systems with Database Management," in W. Reitman (Ed.), Artificial Intelligence Applications for Business, Norwood, NJ: Ablex Publishing Corporation, 1984.

The combined use of database management systems (DBMS) and artificial intelligence-based expert systems (ES) is potentially very valuable for modern business applications. The large body of facts usually required in business information systems can be made available to an ES through an existing commercial DBMS. Furthermore, the DBMS itself can be used more intelligently and operated more efficiently if enhanced with ES features. However, the implementation of DBMS-ES cooperation is not easy. We explore practical benefits of the cooperative use of DBMS and ES, as well as the research challenges it presents. Strategies for providing data from a DBMS to an ES are given. Complementary strategies for providing intelligence from an ES to a DBMS are also presented. Finally, we discuss architectural issues such as degree of coupling, and combination with quantitative methods. As an illustration, a research effort at New York University to integrate a logic-based business ES with a relational DBMS is described.

Johannes, J. D., "Representation and Use of Judgmental Knowledge," Proceedings of the IEEE, New York: IEEE, Nov. 1983.

This paper describes the representation and use of judgmental knowledge in an expert system THYAID. The construction of a judgment-knowledge space and procedures used in developing a thyroid disease diagnostic aid are described. Examples and experimental results are presented as well as the approach used to provide measure of the system's capability.

Johannes, J. D., "Judgmental-Knowledge Bases: Problem Solving and Expert Systems," Proceedings of the International Conference on Data Engineering, Silver Spring, MD: IEEE Computer Society Press, Apr. 1984.

A description is given of the representation and use of judgmental knowledge as applied in the experimental development of a general problem solver and an expert system. Human judgments on the degree to which objects are important to processes (solution methods) are factored, analyzed, and formed into a judgment space of orthogonal vectors. This space is consulted, in the case of the problem solver, to choose a solution method based on the attributes of the available objects. The result is a system that decides how to solve a problem in much the same way a person would. In the expert system THYAID, the attributes of a

patient are combined to form a diagnosis. Examples and experimental results are also given.

Johnson, L., "Inference in the Framework of Expert Systems," IEE Colloquium on Decision Support Aspects of Expert Systems (Digest No. 67), London, England: IEE, May 1984.

The author reports on the forms of inference employed in the reasoning components of the majority of the publicized expert systems. The actual mechanization for implementing these two forms of inference depends on the scheme(s) used for representing the knowledge captured within the systems. These schemes are divided into two categories. One category is the rule-based schemes and the other is a catch-all category (associative nets and frames). How specific systems perform their inference-making is discussed.

Klahr, P., W. S. Faught, and G. R. Martins, "Rule-Oriented Simulation," Proceedings of the International Conference on Cybernetics and Society, New York: IEEE, Oct. 1980.

Large-scale simulators have been plagued with problems of intelligibility, modifiability, credibility and performance. Techniques developed in the area of artificial intelligence, particularly in the development of knowledge-based expert systems, are relevant to the design of simulators that model decision-making behaviors. The authors outline the problems encountered in large-scale simulation and present an approach that incorporates and extends those techniques. ROSS, a rule-oriented simulation system that embodies these techniques, is described.

Klahr, P., D. McArthur, and S. Narain, "SWIRL: An Object-Oriented Air Battle Simulator," Proceedings of the National Conference on Artificial Intelligence, AAAI-82, American Association for Artificial Intelligence, 1982.

Kobler, V. P., "Overview of Tools for Knowledge Base Construction," Proceedings of the International Conference on Data Engineering, Silver Spring, MD: IEEE Computer Society Press, Apr. 1984.

The relative merits of a representative sampling of state-of-the-art tools are described. The tools presented are those that have been publicly documented and are representative of what is available at the present time. They range from efforts to provide a better knowledge representation language to an entire framework with input-output tools.

Kowalski, R., "AI and Software Engineering," Datamation, Nov. 1984.

Kukich, K., "Knowledge-Based Report Generation: A Technique for Automatically Generating Natural Language Reports from Data Bases," Proceedings of the Sixth Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, Baltimore, MD: ACM, June 1983.

Knowledge-based report generation is a technique for automatically generating natural language summaries from data bases. It is so named because it applies the tools of knowledge-based expert systems design to

the problem of text generation. The technique is currently being applied to the design of an automatic natural language stock report generator. Examples drawn from the implementation of the stock report generator are used to describe the components of a knowledge-based report generator.

Laurent, J. P., "Operation of a Knowledge Base: Choices and Strategies," Proceedings of the 4th Congress on Pattern Recognition and Artificial Intelligence, Le Chesnay, France: Inst. Nat. Recherche Inf. and Autom., 1984.

Lee, R. M., "Information System Semantics (A Logic-Based Approach)," Journal of Information Systems, Vol. 1, No. 2, Fall 1984.

Lenat, D. B., "Heuretics: Theoretical and Experimental Study of Heuristic Rules," Proceedings of the National Conference on Artificial Intelligence, AAAI-82, American Association for Artificial Intelligence, 1982.

Builders of expert rule-based systems attribute the impressive performance of their programs to the corpus of knowledge they embody: a large network of facts to provide breadth of scope, and a large array of informal judgmental rules (heuristics) which guide the system toward plausible paths to follow and away from implausible ones. Yet what is the nature of heuristics? What is the source of their power? How do they interrelate; how can/should a large corpus of heuristic rules be organized? How do heuristics originate and evolve? Heuretics is the study of heuristics, with an eye toward answering questions such as those. Two case studies, the AM and EURISKO programs, have led to some tentative heuretics hypotheses, a dozen of which are presented in this paper. Our aim is to stimulate future research in this field.

Lenat, D. B., et al., "Reasoning about Reasoning," in Frederick Hayes-Roth, D. A. Waterman and D. B. Lenat (Eds.), Building Expert Systems, Reading, MA: Addison-Wesley Publishing Company, 1983.

The authors point out that most knowledge-based systems today employ only simple kinds of knowledge in simple ways. They investigate meta-knowledge and meta-expert systems, which aid in the tasks of constructing, maintaining, and extending expert systems. Also, the knowledge engineering and construction methodology issues involved in the building of meta-knowledge and meta-expert systems, especially the need for different kinds of knowledge, are discussed.

Lewis, J. W., "Effective Graphics User Interface for Rules and Inference Mechanisms," Proceedings of the Joint Services Workshop on Artificial Intelligence in Maintenance, Boulder, CO, Oct. 1983.

As the technology of rule-based inference mechanisms matures, knowledge acquisition -- the creation, structuring, and verification of rules -- becomes increasingly important. The accuracy and completeness of the rules in the knowledge base determine expert system performance, and the cost of acquiring that knowledge base dominates all other hardware and software costs in practical systems. To reduce knowledge acquisition time and error rate, a new interactive graphics interface for rules is being designed and implemented in GE Corporate Research and Development. In the new system, each set of rules is represented as an AND/OR graph and parts of the rule base are displayed on a CRT screen as an AND/OR tree. A user -- even an unsophisticated user -- can navigate the AND/OR graph, identify nodes to be modified, analyze the behavior of the graph, verify its correctness graphically, and follow the execution of inference engines.

Lowrance, J. D., and T. D. Garvey, "Evidential Reasoning: A Developing Concept," Proceedings of the International Conference on Cybernetics and Society, New York: IEEE, Oct. 1982.

One common feature of most knowledge-based expert systems is that they must reason based upon evidential information. Yet there is very little agreement on how this should be done. The authors present their current understanding of this problem and some partial solutions. They begin by characterizing evidence as a body of information that is uncertain, incomplete, and sometimes inaccurate. Based on this characterization, the authors conclude that evidential reasoning requires both a method for pooling multiple bodies of evidence to arrive at consensus opinion and some means of drawing the appropriate conclusions from that opinion. This approach, based on a relatively new mathematical theory of evidence, is contrasted with those approaches based on Bayesian probability models. The authors believe that their approach has some significant advantages, particularly its ability to represent and reason from bounded ignorance.

Masui, S., J. McDermott, and A. Sobel, "Decision-Making in Time-Critical Situations," Proceedings of the Eighth IJCAI, American Association for Artificial Intelligence, 1983.

There are a variety of tasks in which the amount of time available to make decisions does not always allow a careful consideration of all of the options. AIRPLAN is being developed to assist with one such task, managing the launch and recovery of aircraft on a carrier. This paper describes attempts to represent AIRPLAN's knowledge in a way that enables it to provide as much assistance as possible in the time available.

McDermott, J., "Extracting Knowledge From Expert Systems," Proceedings of the Eighth IJCAI, American Association for Artificial Intelligence, 1983.

Our understanding of how to use large amounts of knowledge to enhance the problem solving capabilities of computer programs is quite limited. Over the past several years a number of knowledge-based systems have been developed, and this experience has provided us with a handful of techniques we can apply in a few domains. What we don't have yet is much of an appreciation of why these techniques work or of the limits of their usefulness. In order to take this next step, we need more data. Unfortunately, the analyses of the expert systems currently being built tend to ignore questions that could provide precisely the data needed. This paper proposes a few questions that might be worth asking and shows how answers to those questions could begin to give us the understanding we lack.

McDermott, J., and B. Steele, "Extending a Knowledge-Based System to Deal with ad hoc Constraints," Proceedings of the Seventh IJCAI, American Association for Artificial Intelligence, 1981.

R1 is a rule-based program used by Digital Equipment Corporation's Manufacturing Organization for configuring VAX-11 systems. Since its inception, R1 has the capability of producing a functional system by fleshing out an order consisting only of a CPU, some primary memory, and some devices. Until recently, however, it was not capable of accepting as part of its input a set of ad hoc (customer-specific) constraints. Left to itself, R1 was capable of producing reasonable configurations. But it was incapable of modifying its decisions on the basis of information that others could provide. This paper describes how rules were added to R1 to take advantage of such information. R1 can now accept as input commands that specify how particular components are to be configured. Whenever one of the commands becomes relevant, these rules take control, extend the configuration in the direction indicated by the command, and then step aside, allowing R1's ordinary-case configuration rules to regain control.

McDermott, D., and R. Brooks, "ARBY: Diagnosis with Shallow Causal Methods," Proceedings of the National Conference on Artificial Intelligence, AAAI-82, American Association for Artificial Intelligence, 1982.

Michalski, R. S., and R. L. Chilausky, "Knowledge Acquisition by Encoding Expert Rules Versus Computer Induction from Examples: A Case Study Involving Soybean Pathology," International Journal of Man-Machine Studies, Vol. 12, 1980.

In view of the growing interest in the development of knowledge-based computer consulting systems for various problem domains, the problems of knowledge acquisition have special significance. Current methods of knowledge acquisition rely entirely on the direct representation of knowledge of experts, which usually is a very time and effort consuming task. This paper presents results from an experiment to compare the above method of knowledge acquisition with a method based on inductive learning from examples. The comparison was done in the context of developing rules for soybean disease diagnosis and has demonstrated an advantage of the inductively derived rules in performing a testing task (which involved diagnosing a few hundred cases of soybean diseases).

Michie, D., "The State of the Art in Machine Learning," in D. Michie (Ed.), Introductory Readings in Expert Systems, New York: Gordon Breach Science Publishers, 1982.

Moorthy, V. S., and B. Chandrasekaran, "Representation for the Functioning of Devices that Supports Compilation of Expert Problem Solving Structures: An Extended Summary," Proceedings of the Joint Services Workshop on Artificial Intelligence in Maintenance, Boulder, CO., Oct. 1983.

Morris, P. A., "An Axiomatic Approach to Expert Resolution," Management Science, Vol. 29, No. 1, Jan. 1983.

Mylopoulos, J., T. Shibahra, and J. K. Tsotsos, "Building Knowledge-Based Systems: The PSN Experience," Computer, Vol. 16, No. 10, Oct. 1983.

An overview is provided of the procedural semantic networks (PSN) project, whose goal is to develop and test linguistic and other tools that facilitate the construction of large knowledge bases. Two large knowledge-based systems that use PSN as a knowledge-representation language are described. They use the PSN formalism to assess left-ventricular behavior and diagnose cardiac rhythm disorders.

Nau, D. S., J. A. Reggia, and P. Y. Wang, "Knowledge-Based Problem Solving Without Production Rules," *Proceedings of Trends and Applications 1983 on Automating Intelligent Behavior: Applications and Frontiers*, Silver Spring, MD: IEEE Computer Society Press, May 1983.

This paper discusses one alternative to production rules which has proven useful in several medical diagnostic problems and at least one non-medical non-diagnostic problem. This method, which is based on a generalization of the set covering problem, provides an intuitively plausible method for diagnostic reasoning which is sometimes preferable to production rule techniques.

Patil, R. S., P. Szolovits, and W. B. Schwartz, "Information Acquisition in Diagnosis," *Proceedings of the National Conference on Artificial Intelligence*, AAAI-82, American Association for Artificial Intelligence, 1982.

Pau, L. F., "Failure Detection Processes by an Expert System and Hybrid Pattern Recognition," *Pattern Recognition Letters*, Vol. 2, No. 6, Dec. 1984.

Politakis, P., and S. M. Weiss, "A System for Empirical Experimentation With Expert Knowledge," *Proceedings of the Fifteenth Hawaii International Conference on System Sciences*, Vol. 2, Honolulu, HI: Hawaii International Conference of System Sciences, 1983.

Politakis, P., and S. M. Weiss, "Using Empirical Analysis to Refine Expert System Knowledge Bases," *Artificial Intelligence*, Vol. 22, No. 1 Jan. 1984.

SEEK is a system which gives interactive advice about rule refinement during the design of an expert system. The advice takes the form of suggestions for possible experiments in generalizing or specializing rules in a model of reasoning rules cited by the expert. Case experience, in the form of stored cases with known conclusions, is used to interactively guide the expert in refining the rules of a model. This approach is most effective when the model of the expert's knowledge is relatively accurate and small changes in the model may improve performance. The system is interactive; the authors rely on the expert to focus the system on those experiments that appear to be most consistent with his domain knowledge. The design framework of SEEK consists of a tabular format for expressing expert-modeled rules and a general consultation system for applying a model to specific cases. Examples are given from an expert consultation system being developed for diagnosing rheumatic diseases.

Quinlan, J. R., "Discovering Rules by Induction From Large Collections of Examples," in D. Michie (Ed.), *Expert Systems in the Micro-Electronic Age*, Edinburgh, Scotland: Edinburgh University Press, 1979.

Quinlan, J. R., "Semi-Autonomous Acquisition of Pattern-Based Knowledge," in D. Michie (Ed.), Introductory Readings in Expert Systems, New York: Gordon Breach Science Publishers, 1982.

Quinlan, J. R., "Fundamentals of the Knowledge Engineering Problem," in D. Michie (Ed.), Introductory Readings in Expert Systems, New York: Gordon Breach Science Publishers, 1982.

Reboh, R., "Extracting Useful Advice From Conflicting Expertise," Proceedings of the Eighth IJCAI, American Association for Artificial Intelligence, 1983.

A method for automatically identifying areas of disagreement and their sources is presented for multiexpert knowledge-based systems in the context of the Prospector consultation system. It employs performance evaluation techniques in combination with the explanatory facilities present in many expert systems to assist the user of an expert system in deciding which among several possibly conflicting expert opinions he should choose.

Reggia, J. A., D. S. Nau, and P. Y. Wang, "Diagnostic Expert Systems Based on a Set Covering Model," International Journal of Man-Machine Studies, Vol. 19, No. 5, Nov. 1983.

This paper proposes that a generalization of the set covering problem can be used as an intuitively plausible model for diagnostic problem solving. Such a model is potentially useful as a basis for expert systems in that it provides a solution to the difficult problem of multiple simultaneous disorders. The authors briefly introduce the theoretical model and then illustrate its application in diagnostic expert systems. Several challenging issues arise in adopting the set covering model to real-world problems, and these are also discussed along with the solutions they have adopted.

Riesbeck, C., "Knowledge Reorganization and Reasoning Style," International Journal of Man-Machine Studies, Vol. 20, 1984.

Robinson, J. A., "Problems and Trends for the Future of Logic Programming," in D. Michie (Ed.), Introductory Readings in Expert Systems, New York: Gordon Breach Science Publishers, 1982.

Robinson, J. A., "Fundamentals of Machine-Oriented Deductive Logic," in D. Michie (Ed.), Introductory Readings in Expert Systems, New York: Gordon Breach Science Publishers, 1982.

Saurs, R., and R. Walsh, "On Requirements of Future Expert Systems," Proceedings of the Eighth IJCAI, American Association for Artificial Intelligence, 1983.

Shaw, M. L. G., "Acquisition of Personal Knowledge," International Journal of Policy Analysis and Information Systems, Vol. 4, No. 4, Dec. 1980.

Expert systems are increasingly relying on an explicit knowledge base in which the meaning of the stored information must be machine manipulable. This paper presents a method for accessing variables (constructs) which a person is using to categorize conceptional entities, and two interactive computer programs for eliciting these from a human user. The elicitation may commence from an empty knowledge base (PEGASUS), or be directed through an expert's knowledge base (PEGASUS-BANK). The interactive algorithms direct the elicitation of the knowledge base through continuous on-line analysis of the data structures already entered. A detailed example is given of the use of the PEGASUS-BANK program in a medical context.

Slagle, J. R., et al., "An Intelligent Control Strategy for Computer Consultation," IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 6, No. 2, Mar. 1984.

Expert consultant systems often perform computations on a directed graph of associated propositions. Each proposition is represented by a node. Edges connecting these nodes are associated with rules which organize the propositions into antecedent/consequent relationships. A node may be assigned a value through the edges that bind it to its antecedents. Various strategies are employed to determine assignment sequences that result in efficient computer consultation. One such strategy, the merit system, has been successfully implemented in BATTLE, an expert consultant system for the Marine Crops. The merit strategy enables BATTLE to focus the consultation process on the most appropriate questions. The merit system, originally defined for logical functions in the multiple program, has been extended to the MYCIN style of propagation and to the method of subjective Bayesian assignments used by PROSPECTOR. A procedure for merit calculations with any differentiable, real-valued assignment function is presented.

Smith, R. G., "Declarative Task Description as a User-Interface Structuring Mechanism," Computer, Vol. 17, No. 9, Sept. 1984.

Sridharan, N. S., and J. L. Bresina, "Knowledge Structures for Planning in Realistic Domains," Computers & Mathematics, Vol. 11, No. 5, 1985.

Sticklen, J., et al., "MDX-MYCIN: The MDX Paradigm Applied to the MYCIN Domain," Computers & Mathematics, Vol. 11, No. 5, 1985.

Stolfo, S. J., "Knowledge Engineering: Theory and Practice," Proceedings of Trends and Applications 1983 on Automating Intelligence Behavior: Applications and Frontiers, Silver Spring, MD: IEEE Computer Society Press, May 1983.

The authors briefly describe ACE, a knowledge-based system for Automated Automated Cable Expertise, which is designed to provide support for management analysis, automating decision making for telephone cable maintenance. The development of ACE, undertaken by Bell Telephone Laboratories, demonstrates in a forceful way the manner in which AI techniques can be applied to significant and practical real-world problems. ACE is a representative of a new technology, the merging of data base and knowledge-based systems, which may synergistically benefit a large number of applications. However, to fulfill their promise for the very-large-scale embodiment of domain-specific expertise for many real-world problems, systems such as ACE may indeed require an order of magnitude more knowledge. Thus, a number of scientific questions, covering a considerable spectrum of AI research, must be addressed and solved to effect this technology transfer.

Suwa, M., A. C. Scott, and E. H. Shortliffe, "An Approach to Verifying Completeness and Consistency in a Rule-Based Expert System," AI Magazine, Vol. 3, No. 4, 1982.

Swartout, W., "Explaining and Justifying Expert Consulting Programs," Proceedings of the Seventh IJCAI, American Association for Artificial Intelligence, 1981.

Thorndyke, P. W., and K. T. Wescourt, "Modeling Time-Stressed Situation Assessment and Planning for Intelligence Opponent Simulation," Report No.: PPAFTR-1124-84-1, Menlo Park, CA: Knowledge Systems Branch, Perceptronics, Inc., July 1984.

This report summarizes research undertaken to simulate expertise on a real-time time-stressed task requiring situation assessment and planning. While prior efforts in expert systems development have attacked a variety of problems either in situation interpretation or planning, none has attempted to articulate a model that can support both types of activities simultaneously in a dynamic environment. The present study attempted to derive such a cognitive model for time-stressed performance in the domain of Naval surface-warfare tactical decision making. The model extends the concept of the blackboard architecture -- multiple independent processing specialists contributing data and results to a common knowledge base -- previously applied to systems for either signal interpretation or planning. Their approach treats the time-stressed assessment and planning task as a hierarchy of blackboard models, with several processing modules -- each organized as a blackboard model -- interacting through a common knowledge base. Implementation of the model is underway for use as an automated opponent in a Navy tactical battle game. The simulation will emulate an opposing naval commander and will be embedded in an interactive instructional system providing battle exercise for Navy tactics trainees.

Tsotsos, J. K., "Cooperative Computation and Time," Proceedings of the Workshop on Computer Vision: Representation and Control, Silver Spring, MD: IEEE Computer Society Press, 1984.

van Melle, W., "A Domain-Independent Production Rule System for Consultation Programs," Proceedings of the Sixth IJCAI, American Association for Artificial Intelligence, 1979.

EMYCIN is a programming system for writing knowledge-based consultation programs with a production-rule representation of knowledge. The major components of the system, including an explanation program and knowledge acquisition routines, are briefly described. EMYCIN has been used to build consultation systems in several areas of medicine, as well as in the engineering domain. These experiences lead to some general conclusions regarding the potential applicability of EMYCIN to new domains.

Walker, A., "Databases, Expert Systems, and PROLOG," in Reitman (Ed.), Artificial Intelligence Applications for Business, Norwood, NJ: Ablex Publishing Corporation, 1984.

The knowledge needed for an expert system is usually held in the form of facts and IF-THEN rules. The expert system itself is a program (often in LISP) which infers advice from the facts and the rules. In the language PROLOG, a program consists of facts and rules, so the knowledge needed for an expert system can be written down directly. To get advice, one can write further rules about how to use the knowledge. The facts in PROLOG correspond to a relational database, while the rules contain expertise about how to use the data. This chapter uses specific examples, taken from three programs, to show that PROLOG is a practical language for bringing together several apparently diverse techniques in databases and expert systems.

Wall, R. S., "Industrial Strength Knowledge Representation," Proceedings of the Third Annual International Phoenix Conference on Computers and Communications, Silver Spring, MD: IEEE Computer Society Press, 1984.

One of the major problems in trying to build knowledge-based systems is choosing the representation framework that is best suited to the domain. Traditionally, one selects a single paradigm and then encodes the domain knowledge in that framework. This approach is troublesome for large or constantly revised domains. An alternative approach is presented in which a number of different paradigms were implemented as a set of integrated tools. The tools were constructed using the object-oriented programming facility of the LISP machine flavor system. This allows each facet of the domain to be modeled with its most natural representation. The tools include partitioned networks, conceptual hierarchies, inference engines, process simulation and epistemological modeling.

Wall, R. S., and E. L. Rissland, "Scenarios as an Aid to Planning," Proceedings of the National Conference on Artificial Intelligence, AAAI-82, American Association for Artificial Intelligence, 1982.

In many complex domains, full-scale detailed solutions to a problem -- plans and ways to implement them -- are not possible or practical. Rather than attempt to produce such detailed solutions, the authors propose an alternative of using scenarios: a set of "snapshots" of the future that outline possible courses of action, actions of foreign processes, and their consequences. In this paper, they present the scenario idea and discuss its use as an aid to tactical planning. Finally, they discuss an implementation of these ideas in a Tactical Assistant to aid planning in the errand running and conflict simulation game domains and present an example illustrative of its capabilities.

Wallis, J. W., and E. H. Shortliffe, "Explanatory Power for Medical Systems: Studies in the Representation of Causal Relationships for Clinical Consultations," Methods of Information in Medicine, Vol. 21, No. 3, July 1982.

This paper reports on experiments designed to identify and implement mechanisms for enhancing the explanation capabilities of reasoning programs for medical consultation. The goals of an explanation system are discussed, as is the additional knowledge needed to meet these goals in a medical domain. The paper focuses on the generation of explanations that are appropriate for different types of system users. This task requires a knowledge of what is complex and what is important; it is further strengthened by a classification of the associations of causal mechanisms inherent in the inference rules. A causal representation can also be used to aid in refining a comprehensive knowledge base so that the reasoning and explanations are more adequate. The paper describes a prototype system which reasons from causal inference rules and generates explanations that are appropriate for the user.

Waterman, D. A., "A Rule-Based Approach to Knowledge Acquisition for Man-Machine Interface Programs," International Journal of Man-Machine Studies, Vol. 10, 1978.

The development of computer programs, called agents, that act as man-machine interfaces for computer users is described. These programs are written in RITA: the Rule-directed Interactive Transaction Agent system, and are organized as sets of IF-THEN rules or production systems. The programs, or personal computer agents, are divided into two main categories: those that interact the user to computer systems he wishes to use and those that interact with the user to acquire the knowledge needed to create these interface programs. The relationship between the interface program and the knowledge acquisition program is that of parent-offspring. Three types of parent-offspring RITA agent pairs are described: (1) an exemplary programing agent that watches a user perform an arbitrary series of operations on the computer and then writes a program (a task agent) to perform the same task, (2) a tutoring agent that watches an expert demonstrate the use of an interactive computer language or local operating system and then creates a teaching agent that can help naive users become familiar with the language or system demonstrated by the expert, and (3) a reactive-message creating agent which elicits text from a user (the sender) and from it creates a new RITA agent which is a reactive message. The reactive message is sent to some other user (the recipient) who interacts with it. During the course of the interaction a record of the recipient's responses is sent back to the sender.

Weiss, S., et al., "A Model-Based Method for Computer-Aided Medical Decision Making," Artificial Intelligence, Vol. 11, 1978.

A general method of computer assisted medical decision-making has been developed based on causal-associational network (CASNET) models of disease. A CASNET model consists of three main components: observations of a patient, pathophysiological states and disease classifications. As observations are recorded, they are associated with the appropriate states. States are causally related, forming a network that summarizes the mechanisms of disease. Patterns of states in the network are linked to individual disease classifications. Recommendations for broad classes of treatment are triggered by the appropriate diagnostic classes. Strategies of specific treatment selection are guided by the individual pattern of observations and diagnostic conclusions. This approach has been applied in a consultation program for the diagnosis and treatment of the glaucomas.

Yager, R. R., "Querying Knowledge Base Systems with Linguistic Information via Knowledge Trees," International Journal of Man-Machine Studies, Vol. 19, 1983.

The author is interested in finding automated procedures for extracting information from knowledge bases which contain linguistic information. The concept of a possibility distribution to represent the information in a linguistic value is used. The theory of approximate reasoning is used to provide both a means for translating propositions into a machine-understandable form and a methodology for making inferences from this information. An algorithmic procedure resulting in the desired information are presented. This paper is restricted to answering questions about the value of a variable from a knowledge base consisting of simple data statements and implication statements.

Yager, R. R., "Measuring the Quality of Linguistic Forecasts," International Journal of Man-Machine Studies, Vol. 21, 1984

This paper suggests a method for representing linguistic forecasts via fuzzy sets. It uses a representation to obtain a measure of quality of forecasts which takes into account both the validity and the specificity of the forecast.

Yager, R. R., "Approximate Reasoning as a Basis for Rule-Based Expert Systems," Technical Report No. MII-314, Machine Intelligence Institute, Iona College, New Rochelle, NY, 1984.

## 5. TOOLS FOR BUILDING EXPERT SYSTEMS

As the computational and analysis requirements of expert systems are different from those of conventional computer programs, it is not surprising to find that special tools have evolved for the study and development of expert systems, see Barstow, et al., [1983]. Where the needs of conventional programming helped to generate special computer languages like FORTRAN and COBOL, the needs of expert system programming (more accurately, the needs of artificial intelligence) have produced the LISP and PROLOG languages. Similarly, as many conventional problems could be solved more efficiently by standardizing major programming elements into subprograms and subsystems, like file maintenance and retrieval procedures, the design of expert systems can be facilitated by the use of programming subsystems termed shells. Shells are specialty packages that construct the knowledge base and provide an inferencing procedure for exploiting the information contained within the knowledge base. It is also important to note that many tools are now available for building expert systems on microcomputers, see Assad and Golden [1985]. These tools include microversions of LISP and PROLOG, as well as micro-based shells. It is also possible to build the system on a main frame or minicomputer and down load the production version onto a microcomputer. However, once the system is down loaded, it may no longer be possible to modify the knowledge base or control strategy to reflect new information and/or knowledge about the problem.

### 5.1 GENERAL PURPOSE TOOLS: LISP AND PROLOG

LISP and PROLOG provide a general framework for non-numeric processing, especially the heuristics and symbolic manipulation required for knowledge representation. Dating back to 1957, LISP is one of the oldest programming languages still in use. The motivation for the language centered on its list processing capabilities (for which the name LISP is an acronym). LISP has developed along lines that are quite separate from those of the algorithmic languages such as FORTRAN and PASCAL.

LISP interpreters and compilers make the computer behave as if it were a symbol-manipulation machine. Symbols can be used to represent knowledge in terms of objects, relationships that establish meaningful links from one object to another, and processes that govern the creation, destruction, transformation and other behavior of objects. Objects and their relationships can be represented by means of symbolic data structures. An object can have various attributes and their values associated with it by means of a property list. An attribute can stand for both a descriptive property or it can represent a relationship between the described object and another object. Relationships can also be defined in terms of set memberships or predicate functions.

There are only two levels of data objects in LISP. The most basic form is the atom, which is simply a character string used as a symbolic representation. The higher-level data structure is the S-expression (symbolic expression). An S-expression is either a single atom or a linked list of S-expressions. S-expressions are therefore capable of encoding arbitrary tree-structured data. Associated with this linked list data type is an unambiguous one-dimensional surface representation through which the LISP user may communicate. S-expressions are the only form encountered in pure LISP. All programs, as well as all data, conform to this structure and are expressible in the surface representation as parenthesized lists. LISP executes a simple cycle of accepting input S-expressions, evaluating them, and printing out the results of the evaluation. This operation is known as the Read-Eval-Print loop.

The programming language PROLOG was a first attempt to design a language that would enable a programmer to specify tasks in logic, rather than what the machine should do and when. PROLOG stands for "programming in logic." Resolution and unification are formal deductive inference procedures that have been used to prove theorems and to derive information from data that is not explicitly stated, but which is deducible from the stored data. PROLOG implements a goal-driven syntactic strategy based upon the resolution principle. Knowledge in PROLOG is represented in a predicate calculus notation. PROLOG programming differs from LISP, as well as the algorithmic-based languages, in that the programmer uses PROLOG to state the problem, but cannot specify the means or procedure that the system should adopt to achieve a solution.

In PROLOG the control statements are separate from the logic and are invoked in a manner transparent to the programmer. PROLOG also supports arithmetic operators, structures of sets and lists, and their associated manipulations. PROLOG also allows one to define new operators, including their position precedence and associativity.

The semantics of PROLOG is that of resolution logic. All PROLOG statements (axioms) are written in the form "In order to prove B, first prove $A_1$ and $A_2$ and ... and $A_n$." Such a statement is equivalent to the logical formula

$$A_1 \ \& \ A_2 \ \& \cdots \& \ A_n \rightarrow B$$

The A's are referred to as the conditions of the clause and the B's are referred to as the consequences. Knowledge about a problem environment can be stored as axioms in the language. Logical inferences are performed on these axioms to derive new information.

A PROLOG program involves a sentence consisting of a finite set of Horn clauses. Syntactically, a Horn clause is one containing at most one consequent. The method of control used in PROLOG is a restricted form of linear input resolution. Basically, PROLOG adopts a depth-first interpretation of Horn clauses; it considers only one alternative at a time, following up the implications under the assumption that the choice is correct. Programming in PROLOG is not so much like telling a computer what to do when, but rather telling it what is true and asking it to draw conclusions. By looking at a PROLOG program, we would know what it computes rather than how it computes.

## 5.2 SPECIAL PURPOSE TOOLS: SHELLS

Although LISP and PROLOG provide a powerful framework for building expert systems, the recent rapid growth of the field is due to the variety and sophistication of shells. Since the list of shells is expanding, especially for special purpose applications, we shall focus only on five of the better-known systems: (1) ART (Automatic Reasoning Tool), (2) KEE (Knowledge Engineering Environment), (3) OPS5, (4) ROSIE (Rule Oriented System for Implementing Expertise), and (5) TIMM (The Intelligent Machine Model).

ART is a shell tool that provides the system developer with a great deal of representational and programming features, see Williams [1984]. Representational features include: (1) facts, (2) patterns, (3) schemas, (4) rules, and (5) viewpoints. Programming features include: (1) a help system, (2) menus, (3) graphics, and (4) a program monitor to facilitate system debugging during execution.

Propositions in ART express declarative knowledge including general facts and situation specific facts. A fact in ART consists of a proposition (a recursive sequence structure) and an extent (a scope of validity). ART differentiates between facts that are explicitly stated to be false and facts that are unknown (i.e., not known to be true or false). This is supported by allowing ART propositions to have truth values associated with them. In ART, patterns are propositions with variables. Schemas in ART (conceptually similar to frames in KEE) may be thought of as collections of facts, allowing the representation of data as conceptual objects.

Rules are used in ART to encode procedural knowledge. ART provides a rich pattern matching language which is used to create and apply rules. The rule language includes logical connectives and quantification, as well as procedural constraints. Furthermore, mixed chaining is possible by combining goals and facts within a rule. ART differentiates between inference and production rules. Inference rules add facts to the knowledge base, while production rules change facts.

Viewpoints are an interesting feature provided by ART. A viewpoint is a hypothetical world in which ART can pursue a potential course of action to see where it leads. Hypotheticals may be compared and contrasted, merged together or marked as inconsistent, automatically or by rules. Viewpoints can be used to represent a point in time, a possible interruption of incoming data, or a potential solution to a problem. Associated with each viewpoint is a confidence rating that gives a numeric estimate of a viewpoint's validity. Numerical certainty factors can also be associated with individual facts in the knowledge base.

KEE is a shell that integrates several artificial intelligence methodologies into a single system, see Intelli Corp [1984] and Finnucan [1985]. It provides support for frame-based knowledge representation with taxonomic inheritance, rule-based reasoning, logic representations, data-driven reasoning, object-oriented programming, and LISP functional programming. In addition, KEE provides direct support for various uses of interactive graphics. Important ideas in KEE include early prototype development, incremental refinement of the problem description, use of multiple integrated solution methods, and emphasis on visibility of both the problem-solution process and the explicit description of the problem domain.

Application specific knowledge in KEE is called a knowledge base and is loaded into the KEE system. The knowledge base includes frames that represent or describe objects and classes of objects. Frame systems are inherently object-oriented as the frames each represent an object (a class is also an object). Behavior is invoked by passing messages to objects. The information needed to respond to a message is stored in the attribute slots. Storing behavioral responses directly with an object can facilitate creating more modular programs.

Frames in KEE are a data structure consisting of a frame name and a set of attribute descriptions, each defined by a slot. Associated with each slot is the attribute's value(s), constraints that the value must satisfy, procedures that are called whenever the slot is accessed or its values change, and the source of the inherited values. Attribute values can be restricted to meet constraints including class, range, cardinality, and data type. Objects include classes, subclasses, and individuals.

Rule clauses in KEE can include logical expressions with conjunctions, disjunctions, and negations, as well as arbitrary LISP S-expressions. KEE also supports backward and forward chaining rule interpreters. The system permits rules to be used for multiple purposes such as explanation by describing or displaying a chain of reasoning, as the chain of rules are applied. User-supplied functions can tailor rule interpretation to specify depth-first, breadth-first, or other search strategies for backward chaining, and resolve conflicts in forward chaining when more than one rule is applicable.

OPS5 is a rule-based programming language descended from earlier OPS languages designed for artificial intelligence and cognitive psychology applications, see Waterman and Hayes-Roth [1983]. Data elements in OPS5 are either vectors or objects with associated attribute-value pairs. Rules in OPS5 are data-driven and operate on a single global data base. The rules have an antecedent, a partial description of data elements, and a consequent, one or more action to be taken if the antecedent matches the data base.

Control in OPS5 is governed by the recognize-act cycle, a simple loop in which rules with satisfied antecedents are found, one is selected (conflict resolution), and its actions are performed. Most OPS5 programs make use of goals to direct the processing, putting them into the data base and deleting them when appropriate for the current problem-solving strategy. The collection of rules that are sensitive to the goal for a given task compose the "method" for that task. Methods and goals are not features of the OPS5 language; they are simply convenient organizing principles. The rules associated with a method can perform tasks directly, or they can create goals that ask other rules to perform the tasks.

ROSIE is a general-purpose rule-based programming system suitable for a broad range of knowledge engineering applications, see Fain, et al., [1982]. Its most striking feature is its English-like syntax, which facilitates the creation and manipulation of the ROSIE data base. This data base contains general relationships with English correspondences in the ROSIE language. To test relations in the data base, the user writes conditions in English that correspond to the relational forms. Five basic types of English relationships are modeled by the language: (1) class membership, (2) predication, (3) intransitive verbs, (4) transitive verbs, and (5) predicate complements. These relationships can be further refined by the use of prepositional phrases.

ROSIE rules are organized as rule sets, defined to be either procedures, generators, or predicates. Generators return values of any sort, while predicates determine proposition truth values. The rule sets operate on a global data base, but maintain temporary private data bases during execution.

ROSIE supports three types of inference mechanisms: (1) state-driven, where the state of the system directly causes a rule to fire, (2) goal-driven, where backward-chaining is used to find rules that will verify predicates in the rule conditions, and (3) change-driven, where a data-base change causes a rule to fire.

TIMM is a general-purpose rule-based programming system, see Craig [1986] and McWhite [1985]. TIMM is an interactive tool that performs the function of the knowledge engineer by querying the expert, and constructing and interpreting the knowledge base. This is accomplished by applying a general method for building an expert system and solving problems that are independent of the problem type.

TIMM's principal functions are: (1) build, (2) train, and (3) exercise. The build function enables one to define the decision to be made; possible decision options (choices), factors affecting the decision, and possible factor values. The train function is used to define rules, to identify generalities among the rules, to compress the knowledge base to its most efficient and effective size, and to identify inconsistencies in the knowledge base. The exercise function is used to test the effectiveness of the expert system in reaching an appropriate decision.

TIMM reasons by analogical or pattern-directed inference, comparing the current situation to similar experiences (rules) in its knowledge base. When an exact match occurs between the antecedent clauses of a rule and the situation, TIMM's decision is the related consequent clause of the rule. If there is no exact rule match, TIMM will still provide a decision. For each rule in the knowledge base, TIMM computes a numerical distance between the rule and the current situation. If the distance equals zero, then an exact match has occurred; if the distance is a non-zero number, the distance provides a measure of how well the current situation is matched by the rule. This is reflected as a reliability measure. Because of this capability, the knowledge base can be searched for near matches in response to precise and imprecise questions.

The TIMM software is written in FORTRAN 77. This approach should promote transportability of the system. In addition, TIMM can be called by another FORTRAN program.

## References

Anderson, R. H., and J. J. Gillogly, "The RAND Intelligence Terminal (RITA) as a Network Access Aid," Proceedings of the 1976 AFIPS National Computer Conference, Volume 45, Arlington, VA: AFIPS Press, 1976.

Artem'eva, I. L., et al., "Compiler Generator for Knowledge Respresentation Languages," Transactions in Programming and Computer Software, Volume 9, No. 4, 1983.

This article describes a compiler generator for automatic realization of expert systems. The specific features of the realization process are described, and a framework is developed for compiling a knowledge base into an expert system, distinct from the traditional method of knowledge base interpretation.

Assad, A. A., and B. L. Golden, "Expert Systems, Microcomputers and Operations Research," Report No.: MS/85-025, College of Business and Management, University of Maryland, College Park, MD, July 1985.

The goal of this article is to review microcomputer-based software for building expert systems. There are now at least 15 commercially-available tools to facilitate the development of expert systems on microcomputers. First, the field of expert systems is introduced in general terms. Then, issues relating more directly to microcomputers are studied in detail. Trends and future directions are also discussed.

Balzer, R., et al., "Hearsay-III: A Domain-Independent Framework for Expert Systems," Proceedings of the National Conference on Artificial Intelligence, AAAI-80, American Association for Artificial Intelligence, 1980.

Barstow, D. R., et al., "Languages and Tools for Knowledge Engineering," in F. Hayes-Roth, D. A. Waterman, and D. B. Lenat (Eds.), Building Expert Systems, Reading, MA: Addison-Wesley Publishing Company, 1983.

In this article, the authors provide a comprehensive survey of tools widely used in knowledge engineering. Each tool reflects a special emphasis on kinds of problems, kinds of overall designs, or kinds of implementations deemed desirable. Although none of the tools meets industrial standards for wide-scale utilization, an order-of-magnitude reduction in development time can be realized by applying an appropriate tool. Today's tools foreshadow a vast array of future industrial implements. Knowledge engineering means forming and assembling knowledge, and the tools make this possible.

Brachman, R. J., and H. J. Levesque, "Competence in Knowledge Representation," Proceedings of the National Conference on Artificial Intelligence, AAAI-82, Amercian Association for Artificial Intelligence, 1982.

The range of domains and tasks for "knowledge-based systems" has been expanding at a furious pace. As we move away from trivial domains, such as the blocks world, the demands on knowledge representation systems used by expert programs are becoming more extreme. For one thing, the domains themselves are getting so complex that specialized technical vocabularies are unavoidable; consequently, the issue of a system talking with an expert in his own language cannot be ignored. For another, tasks such as medical diagnosis, scene analysis, speech understanding, and game playing all have as a central feature an incrementally evolving model representing probably incomplete knowledge of part of the task domain. In this paper, we explore some of the impact of these two critical issues -- complexity and incompleteness -- on knowledge representation systems. We review some aspects of current representation research that offer a foundation for coping with these problems, and finally suggest a way of integrating these ideas into a powerful, practical knowledge representation paradigm.

Brooks, R., "A Comparison Among Four Packages for Knowledge-Based Systems," Proceedings of the International Conference on Cybernetics and Society, New York: IEEE, Oct. 1981.

Four tools for the construction of knowledge-based systems, EMYCIN, EXPERT, AMIDS, and the blackboard option in AGE, are compared in regard to their capabilities for representing three kinds of knowledge: judgmental or decision-making knowledge, explanatory knowledge, and modeling knowledge.

Buchanan, B. G., et al., "Constructing an Expert System," in F. Hayes-Roth, D. A. Waterman, and D. B. Lenat (Eds.), Building Expert Systems, Reading, MA: Addison-Wesley Publishing Company, 1983.

Bylander, T., S. Mittal, and B. Chandrasekaran, "CRSL: A Language for Expert Systems for Diagnosis," Computers and Mathematics, Volume 11, No. 5, 1985.

Charniak, E., C. K. Riesbeck, and D. V. McDermott, Artificial Intelligence Programming, Hillsdale, NJ: Lawrence Erlbaum Associates, Publishers, 1980.

Clark, K. L., and F. G. McCabe, "PROLOG: A Language for Implementing Expert Systems," in J. E. Hayes, D. Michie, and Y. H. Pao (Eds.), Machine Intelligence 10, Chichester, England: Horwood, 1982.

Craig, J. N., "A Tool for Building Expert Systems," Proceedings of the First Symposium on the Theory and Application of Expert Systems in Emergency Management Operations, Special Publication 717, Gaithersburg, MD: National Bureau of Standards, 1986.

TIMM is a flexible and practical system for developing and using expert systems. TIMM allows experts to interact directly with the program in constructing knowledge bases by automating the task of knowledge acquisition. TIMM interviews the expert to define the domain, poses situations for the expert to consider, induces general rules from specific cases, and can include expertise from multiple experts. Rapid prototyping and iterative development of expert systems are supported by TIMM. The basic TIMM functions can be viewed as a toolbox and each of the tools can be used to perform a task during system development. The ability to modify and refine decision domains and rules allows basic systems to be constructed quickly. By linking multiple expert systems together, both structure and hierarchy can be improved on the overall system design. TIMM is written in Fortran and runs in standard computing environments ranging from IBM PC class computers to large mainframes. No special programming is required to input expert systems using TIMM. Thus, the critical area of knowledge acquisition can be addressed directly.

Duffin, P. and S. F. Lello, "MASES - A PROLOG Expert System," Information Age, Volume 6, No. 4, 1984.

The development of an expert system using PROLOG is described. The system, Microcomputer Advice and Selection Expert System (MASES), was intended to provide advice on the selection of microcomputers. However, unforeseen limitations of the microcomputer environment in which MASES was developed prevented achievement of the initial aim of a live, usable system, but a small system, capable of exhibiting some typical features of expert systems, was devised for use in demonstrations. The construction demonstrates how additions to or deletions from the knowledge base can be made without consequent amendment to the set of operating rules. The working model is limited to three microcomputers (Rainbow, Torch and Videcom) and contains 29 rules and 65 facts about printers, interfaces, software and protocols on four linked knowledge bases. It enables users' requirements to be matched against the facilities available on the three microcomputers.

Ennis, S. P., "Expert Systems: A User's Perspective of Some Current Tools," Proceedings of the National Conference on Artificial Intelligence, AAAI-82, American Association for Artificial Intelligence, 1982.

The purpose of this paper is to report on one user's experience with several of the software tools for building an expert system. To the best of the author's knowledge, this is the first time a number of different expert system building tools have been applied to a single problem by a single analyst.

Erman, L. D. and V. R. Lesser, "Hearsay-II: Tutorial Introduction and Retrospective View," Report No.: CMU-CS-78-117; AFOSR-TR-78-1246, Department of Computer Science, Carnegie-Mellon University, Pittsburgh, PA, May 1978.

The Hearsay-II system, developed at CMU as part of the five-year ARPA speech-understanding project, was successfully demonstrated at the end of that project in September 1976. This report reprints two Hearsay II papers which describe and discuss that version of the system: 'The Hearsay-II System: A Tutorial', and 'A Retrospective View of the Hearsay-II Architecture'. The first paper presents a short introduction to the general Hearsay-II structure and describes the September 1976 configuration of knowledge-sources; it includes a detailed description of an utterance being recognized. The second paper discusses the general Hearsay-II architecture and some of the crucial problems encountered in applying that architecture to the problem of speech understanding.

Erman, L. D., P. E. London, and S. F. Fickas, "The Design and an Example Use of HEARSAY-III," Proceedings of the Seventh IJCAI, American Association for Artificial Intelligence, 1981.

HEARSAY-III provides a framework for constructing knowledge-based expert systems. While HEARSAY-III makes no commitment to any particular applicable facilities. These include representation primitives and an interpreter for large-grained, flexible schedulable production rules called knowledge sources. A detailed overview of the motivations behind HEARSAY-III and the facilities it provides are presented. Finally, an application of HEARSAY-III is described.

Fahlham, S. E., NETL: A System for Representing and Using Real-World Knowledge, Cambridge, MA: MIT Press, 1979.

Fain, J., et al., "The ROSIE Language Reference Manual," Report No.: RAND/N-1647-ARPA, RAND Corp., Santa Monica, CA, Dec. 1981.

ROSIE is a programming system designed to support the development of expert systems and other heuristic programming applications. ROSIE offers great promise because it combines many modern capabilities within a single system, including English as a programming language; knowledge representation capabilities; a built-in relational databases system; flexible string matching and communication capabilities; highly modular programming structures; a file package and interactive programming environment. This note describes in detail the features of the ROSIE environment, and the syntax and semantics of the ROSIE language. The Reference Manual should be useful to readers who intend to program in ROSIE or who are interested in the implementational details of the ROSIE system.

Fain, J., et al., "Programming in ROSIE: An Introduction by Means of Examples," Report No.: RAND/N-1646-ARPA, Rand Corp., Santa Monica, CA, Feb. 1982.

ROSIE is a programming system designed to support the development of expert systems and other heuristic programming applications. ROSIE offers great promise because it combines many modern capabilities within a single system, including English as a programming language; knowledge representation capabilities; a built-in relational database system; flexible string matching and communication capabilities; highly modular programming structures; a file package and interactive programming environment. This note presents a number of examples of ROSIE programs in such diverse areas as legal decision making, war gaming, and detecting hazardous chemical spills. The note should be useful to readers who intend to program in ROSIE.

Feigenbaum, E. A., "Knowledge Engineering: The Applied Side," in J. E. Hayes and D. Michie (Eds.), Intelligent Systems: The Unprecedented Opportunity, Chichester, England: Ellis Horwood Limited, 1983.

This paper discusses the applied artificial intelligence work that is sometimes called knowledge engineering. The work is based on computer programs which do symbolic manipulations and symbolic inference, not calculation. The programs discussed do essentially no numerical calculation. They discover qualitative lines-of-reasoning leading to solutions to problems stated symbolically.

Georgeff, M., and U. Bonollo, "Procedural Expert Systems," Proceedings of the Eighth IJCAI, American Association for Artificial Intelligence, 1983.

Glasgow, J., and R. Browse, "Programming Languages for Artificial Intelligence," Computers and Mathematics, Volume 11, No. 5, 1985.

Greiner, R., "RLL-1: A Representation Language Language," Report No.: HPP-80-9, Department of Computer Science, Stanford University, CA, Oct. 1980.

The field of AI is strewn with knowledge representation languages. The language designer typically designs that language with one particular application domain in mind; as subsequent types of applications are tried, what had originally been useful features are found to be undersirable limitations, and the language is overhauled or scrapped. One remedy to this bleak cycle might be to construct a representation language whose domain is the field of representational languages itself. Toward this end, the authors have designed and implemented RLL-1, a frame-based Representation Language Language. The components of representation languages in general (such as slots and inheritance mechanisms) and of RLL-1 itself, in particular, are encoded declaratively as frames. By modifying these frames, the user can change the semantics of RLL-1's components, and significantly alter the overall character of the RLL-1 environment.

Greiner, R., and D. B. Lenat, "RLL-1: A Representation Language Language, Supplement, Details of RLL-1," Report No.: HPP-80-23, Department of Computer Science, Stanford University, CA, Oct. 1980.

This paper includes many implementation level details about the RLL-1 system, described in a companion paper, RLL-1: A Representation Language Language. The contents are as follows: Special Units, Naming conventions, Legend, Actual Units, Index of Units, Environment, Top Level Functions, Functions Needed to Bootstrap RLL-1, Convenience Functions, Advised Functions, and Global Variables.

Gyanfas, F., and M. Popper, "CODEX: Prototypes Driven Backward and Forward Chaining Computer-Based Diagnostic Expert System," Proceedings of the Seventh European Meeting on Cybernetics and Systems Research 2, Amsterdam, Netherlands: North-Holland, 1984.

Introduces a medical aid computer-based diagnostic expert system, CODEX. As it is an empty expert systems requiring an attached knowledge base, its use is not limited to medicine only. The main principles of the system's knowledge base design and the system performance are described. Some implementation details and information on application domains are given.

Hadley, R. F., "SHADOW: A Natural Language Query Analyzer," Computers and Mathematics, Volume 11, No. 5, 1985.

Harrison, N., "Knowledge Base Builders," Systems International, Volume 12, No. 8, Aug. 1984.

Hayes-Roth, F., et al., "Rationale and Motivation for ROSIE," Report No.: RAND/N-1648-ARPA, RAND Corp., Santa Monica, CA, Nov. 1981.

ROSIE (Rule-Oriented System for Implementing Expertise) is a programming language and programming system for artificial intelligence (AI) applications. The ROSIE language is a stylized version of English. The primary design goal for the language was to achieve exceptional program readability; a second goal was to support the development of significant applications. ROSIE allows the programmer to describe complex relationships simply and to manipulate them symbolically and deductively. In addition, it supports network communications and patterned reading and writing to other systems. It also provides for interactive, compiled, and interpreted computing, with a variety of debugging and programming tools. ROSIE encompasses many of the capabilities of conventional programming languages. It is a general-purpose language offering a variety of typical data types and control constructs found in most high-level languages, together with a few found only in AI languages. Among the features included are: rulesets that generate sets, predicates that test propositions, propositional data types, and some limited forms of deduction. Using ROSIE, an AI programmer can think concretely about the problem domain and translate ideas into a program using substantially the same vocabulary that arises in the English (non-computational) formulation of the model.

Intellicorp, "The Knowledge Engineering Environment," Knowledge Systems Division, Mountain View, CA, 1984.

Johnson, W. L., and E. Soloway, "Prost: Knowledge-Based Program Understanding," Proceedings of the Seventh International Conference of Software Engineering, New York: IEEE, Mar. 1984.

A program called Proust is described which does on-line analysis and understanding of PASCAL programs written by novice programmers. Proust takes as input a program and a nonalgorithmic description of the program requirements, and finds the most likely mapping between the requirements and the code. This mapping is in essence a reconstruction of the design and implementation steps that the programmer went through in writing the program. A knowledge base of programming plans and strategies, together with common bugs associated with them, is used in constructing this mapping. Bugs are discovered in the process of relating plans to the code: Proust can therefore give deep explanations of program bugs by relating the buggy implementation steps that the programmer went through in writing the program. A knowledge base of programming plans and strategies, together with common bugs associated with them, is used in constructing this mapping.

Kinnucan, P., "Software Tools Speed Expert System Development," High Technology, Mar. 1985.

Konopasek, M., and S. Jayaraman, "Expert Systems for Personal Computers - The TK! Solver Approach," Byte, Volume 9, No. 5, May 1984.

Although some question the possiblity of implementing expert systems on today's personal computers, the authors of this article wish to present a counter-example. This article shows that many characteristics of expert systems are present in TK! Solver and describes a framework for building a variety of expert systems with quantifiable knowledge bases. TK! Solver is aimed at realizing many concepts expounded in AI research, human-computer interface design, and human problem solving. It has no built-in knowledge of any particular discipline, but it provides a framework to make it easier for the user with such knowledge to construct expert systems. The knowledge engineer - the bottleneck - is eliminated. TK! Solver was designed to be an expert system primarily in the area of numerical problem solving. As such it (1) passes entered algebraic equations and generates a list of variables, (2) solves sets of equations using the consecutive-substitution procedure, (3) solves sets of simultaneous algebraic equations by a modified Newton-Raphson iterative procedure when the consecutive-substitution procedure fails, (4) searches through tables of data and evaluates either unknown functions, values or arguments when required in the process of (2) or (3), (5) performs unit conversions, (6) detects inconsistencies in problem formulation and domain errors, (7) generates a series of solutions for lists of input data and outputs results tabular and graphic forms.

Kukich, K., "Design of a Knowledge-Based Report Generator," Proceedings of the Twenty-first Annual Meeting of the Association for Computational Linguistics, Menlo Park, CA: Association of Computational Linguistics, June 1983.

Leal, A., "Evaluating the Effectiveness of Military Support Systems: Theoretical Foundations, Expert System Design, and Experimental Plan (Research note)," Report No.: ISC-345-3; ARI-RN-83-18, Integrated Sciences Corp., Santa Monica, CA, Sept. 1982.

The main objective of this program is to construct a flexible testbed for the evaluation of the effectiveness of computer-based expert systems in military training and planning. The technical approach consists of simulating the characteristics of expert systems in a game-like environment. Such characteristics include friendly system user interaction, system explanations of rationale about decision recommendations, an ability to make relevant suggestions and comments about situation assessments and about plans proposed by the user, and the use of high-level strategic concepts and terminology. The required software for such a program includes a game environment simulator called the Scenario Generator, a simulated expert system for the game called the Expert Aid, an Optimality Algorithm for computing the best decisions in any situation, and an Evaluation Module for recording execution histories and performance parameters. The expert system will monitor the progress of the game and can be interrogated as the user sees fit. A facility will also be provided for evaluating the user's performance under different modes of consultation with the expert system.

Lenat, D. B., "Computer Software for Intelligent Systems," Scientific American, Sept. 1984.

McDermott, J., "R1: An Expert in the Computer Systems Domain," Proceedings of the National Conference on Artificial Intelligence, AAAI-80, American Association for Artificial Intelligence, 1980.

McDermott, J., "R1: The Formative Years," AI Magazine, Vol. 2, 1981.

McDermott, J., "Building Expert Systems," in W. Reitman (Ed.), Artificial Intelligence Applications for Business, Norwood, NJ: Ablex Publishing Corporation, 1984.

Though only a handful of expert systems have as yet found a home in industry, patterns in their development histories are beginning to emerge. This chapter describes four stages that developers of expert systems typically pass through and then exemplifies these stages with three case studies.

McKeown, Jr., D. M., and J. McDermott, "Toward Expert Systems for Photo Interpretation," Proceedings of Trends and Applications 1983 on Automating Intelligent Behavior: Applications and Frontiers, Silver Spring, MD: IEEE Computer Society Press, May 1983.

McWhite, P., "TIMM: A Fortran-Based Expert Systems Applications Generator," General Research Corporation, McLean, VA, Feb. 1985.

This paper presents TIMM capabilities as used in the development of an actual expert system. Examples demonstrate functions which recommend new rules to make the knowledge base more complete, check for consistency, and generalize knowledge. Notably, the generalize function helps find the key factors affecting a decision. TIMM does not require special compilers, computers, programmers, or personnel. It can be used immediately on most popular computers and operated by experts with minimal training. As a result, it offers a cost-effective alternative to other methodologies for building and implementing expert systems. TIMM supports a pattern directed inference scheme. TIMM can define, modify, exercise, and link expert systems. It accepts uncertain data, makes decisions in uncertain situations, and provides a reliability measure for each decision. No computer programming is required and most of the knowledge engineering function is automated so that subject matter experts can directly develop knowledge bases. Advantages over other methods of developing expert systems, especially those using LISP, include transportability, reduced memory requirements, and rapid execution.

Merry, M., "APEX 3: An Expert System Shell for Fault Diagnosis," GEC Journal Research Incorporated Marconi Review, Volume 1, No. 1, 1983.

An expert system is a computer program designed to simulate the behavior of a human expert in a given field. It consists of a knowledge base of expert knowledge about the field, and an inference engine, to make deductions from this knowledge. An expert system shell is an empty expert system; an inference engine requiring a knowledge base to make an actual expert system, but able to be used with several different knowledge bases to make several different expert systems. APEX 3 is an expert system shell for building fault diagnosis expert systems. APEX 3 is described in detail, paying particular attention to its control strategies, which are designed to enable the system to make diagnoses by asking the user as few questions as possible. The author describes the commands available to the user, and gives some notes on the implementation of APEX 3.

Mizoguchi, F., "A Software Environment for Developing Knowledge Base Systems," in T. Kitagawa (Ed.), Computer Science and Technologies, Amsterdam, Netherlands: North-Holland, 1982.

This paper describes a software environment for developing knowledge-based systems. The system called Multi-Layered Software Environment (MLSE) is proposed for providing a designer of an expert system with wide varieties of design alternatives in software components which derive from artificial intelligence (AI) technolgoy. The MLSE is a collection of module packages for building the tools for knowledge-based systems. That is, by integrating the AI Technology into the MLSE scheme, the author has emphasized the layered approach to build the software environment as the basis for developing knowledge-based systems. To illustrate the strategy of the MLSE scheme the three case studies for designing production systems are described. Based upon these case studies, a new production system called ADIPS was developed which was a product of the MLSE scheme. The MLSE has been implemented in LISP F-3 on M200 of HITAC, IBM 3031 and Inter-LISP of DEC-20.

Mizoguchi, F., "PROLOG Based Expert System," New Generation Computers, Volume 1, No. 1, 1983.

A PROLOG based expert system called APLICOT is described. Its basic design framework is the same as used for other languages implementing expert systems. This utilization of the same design framework permits a comparison to be made of system performance obtained within the same domain of knowledge. The domain chosen involves a reactor's fault diagnostic system consisting of 76 production rules. The results of this comparison show the overall performance of APLICOT to be of the same level as that of EXPERT, EMYCIN and ADIPS, which were developed in the past utilizing different programming languages. Although its code size is ten times smaller than that of LISP or FORTRAN based expert systems, APLICOT's backward and forward reasoning system gives it the same level of system performance and flexible inference strategy as these other systems. It thus demonstrates the potential software productivity of PROLOG-based expert systems.

Morris, P. A., "An Axiomatic Approach to Expert Resolution," Management Science, Vol. 29, No. 1, Jan. 1983.

This paper addresses the problem of resolving the differing probability judgments of a group of experts. Such judgments may be probability assignments if the uncertainty of interest is on the outcome of an event such as the occurrence of rain or the winner of an election. Or, the judgments may be probability densities or distribution functions if the uncertainty is on a continuous factor, such as the length of the Golden Gate Bridge or the market share of a new product. The axiomatic approach proposed here provides a unified Bayesian theory for addressing both types of problems, and provides a framework for understanding why some common approaches to the problem are inadequate.

Naylor, C., "ES/P Advisor," Practical Computer, Vol. 7, No. 10, Oct. 1984.

The author reviews ES/P Advisor, an expert system which enables users to produce a knowledge base, which can then be run with a standard inference engine. Its prime market is expected to be firms who want to automate standard procedures and regulations for use by unskilled staff. ES/P Advisor is available for machines running CP/M-86, MS-DOS and PC-DOS.

Neiman, D., "Graphical Animation from Knowledge," Proceedings of the National Conference on Artificial Intelligence, AAAI-82, American Association for Artificial Intelligence, 1982.

Newell, A., "HARPY, Production Systems and Human Cognition," Report No.: CMU-CS-78-140, Computer Science Department, Carnegie-Mellon University, 1978.

Oliveira, E., "Developing Expert Systems Builders in Logic Programming," New Generation Computers, Vol. 2, No. 2, 1984.

The author develops a set of kits to build expert systems using PROLOG. Two principal modules, a knowledge base acquisition and consultation subsystems are presented. Several knowledge representation structures and mixed inference mechanisms are proposed for the sake of system efficiency. Finally, some explanation capabilities derived accordingly with used inference methods are also implemented and presented.

Parker, R. E., and S. J. Kiselewich, "The Modeling of Human Cognitive Decision Processes in The Intelligent Machine Model (TIMM)," Proceedings of the 1984 American Control Conference, Vol. 1, New York: IEEE, June 1984.

The Intelligent Machine Model (TIMM) is a software package that enables a user to build an expert system, that is, a system which is capable of providing expert advice in some well-defined domain of expertise. To accomplish this, TIMM attempts to represent some of the analog character of the real-world domain, in addition to a standard rule base. This extra information defines a metric over the rule base. This metric allows TIMM to model some distinctly human capabilities. TIMM is able to reach decisions for new, uncertain, and incompletely defined situations. TIMM is able to examine its own rule base and suggest new rules that seem appropriate for the domain. These are important capabilities for real-world expert systems.

People, Jr., H. E., "Knowledge-Based Expert Systems: The Buy or Build Decision," in W. Reitman (Ed.), Artificial Intelligence Applications for Business, Norwood, NJ: Ablex Publishing Corporation, 1984.

Reboh, R., "Knowledge Engineering Techniques and Tools in the PROSPECTOR Environment," Technical Report 8172, SRI International, Menlo Park, CA, June 1981.

Rich, C., "Knowledge Representation Languages and Predicate Calculus: How to Have Your Cake and Eat it Too," Proceedings of the National Conference on Artificial Intelligence, AAAI-82, American Association for Artificial Intelligence, 1982.

This paper attempts to resolve some of the controversy between advocates of predicate calculus and users of other knowledge representation languages by demonstrating that it is possible to have the key features of both in a hybrid system. An example is given of a recently implemented hybrid system in which a specialized planning language co-exists with its translation into predicate calculus. In this system, various kinds of reasoning required for a program understanding task are implemented at either the predicate calculus level or the planning language level, depending on which is more natural.

Richer, M. H., "Evaluating the Existing Tools for Developing Knowledge-Based Systems," Report No. KSL85-19, Stanford Knowledge Systems Laboratory, Stanford University, CA, May 1985.

The purpose of this paper is two-fold: (1) to examine the criteria for evaluating software tools for developing KB systems, and (2) to describe several existing KB system tools. The reader should take note that the information gathered here is necessarily limited to a large extent by what is publicly available. A more complete evaluation must include actual use of each system, and this was not possible at the time this paper was prepared. Therefore, no attempt is made to present a complete evaluation of any one system, and strong comparisons about two or more systems are avoided. Because this paper is not meant to be tutorial in nature, no attempt is made to systematically explain terms and concepts that are commonly used in the artificial intelligence community.

Sheil, B., "The Artificial Intelligence Tool Box," in W. Reitman (Ed.), Artificial Intelligence Applications for Business, Norwood, NJ: Ablex Publishing Corporation, 1984.

Artificial intelligence is often claimed to be applicable to a variety of commercial applications. The author argues that a significant part of that applicability derives, not from the AI functionality per se, but from the underlying software technology. This software technology is examined, its genesis in the nature of AI programming is described, and the ways in which it differs from conventional software engineering are explained. The effectiveness of the AI programming technology for conventional application development first became clear when it was used to construct the information management environments in which the first AI applications were packaged. Building such environments is known to be extremely difficult using conventional software technology. But, in many ways, AI programming technology turns out to be ideally suited to this kind of development. The ability to build highly customized information management systems is itself an application opportunity of enormous leverage. Thus, in the short run at least, the application development power of AI's software technology may turn out to be just as important as the AI techniques for which it is nominally just a carrier.

Stefik, M., et al., "The Organization of Expert Systems: A Prescriptive Tutorial," Xerox Palo Alto Research Centers, Palo Alto, CA, 1982.

Problems in reasoning are diverse. Solving them requires substantial and varied capabilities of problem solvers. This tutorial is about the design of expert systems, that is, systems for solving problems that require substantial expertise. It focuses on principles for matching the organization of an expert system to the difficulties posed by its problems. The tutorial begins with a brief review of standard topics from artificial intelligence, such as the use of the predicate calculus and search techniques. It also discusses some newer topics such as the use of abstractions, assumptions, dependency information and meta-level cognition. These topics are essential for understanding current work on expert systems. The main section of the tutorial is a discussion of organizational prescriptions for problem solvers. It begins with a restricted class of problems that admits a very simple organization. This organization is appropriate when the input data and knowledge are static and reliable and the solution space is small enough to search exhaustively. Few real world problems satisfy all of these requirements. The requirements are then relaxed, one at a time, yielding ten case studies of organizational prescriptions. The first cases describe techniques for dealing with unreliable data and time-varying data. Other cases illustrate techniques for creating abstract solution spaces and using multiple lines for reasoning. The prescriptions are compared for their coverage and illustrated by examples from recent expert systems.

Stefik, M., et al., "Basic Concepts for Building Expert Systems," in F. Hayes-Roth, D. A. Waterman, and D. B. Lenat (Eds.), Building Expert Systems, Reading, MA: Addison-Wesley Publishing Company, 1983.

The authors of this chapter provide, in a summary form, the concepts involved in the general field of artificial intelligence, and several ideas that are central to knowledge engineering. Knowledge typically derives its value from its potential to pare down to manageable size problems that appear impossibly complex. This power comes from its capacity to shrink large search spaces. Search techniques that cannot apply knowledge will not succeed at difficult tasks. A description of many new types of search structures and techniques utilizing knowledge that have evolved is given. And finally, the usefulness of these in a large number of potential applications that make use of similar approaches is provided.

Stefik, M., et al., "The Architecture of Expert Systems," in F. Hayes-Roth, D. A. Waterman, and D. B. Lenat (Eds.), Building Expert Systems, Reading, MA: Addison-Wesley Publishing Company, 1983.

Surya, "Expert-Ease," Personal Computer World, Volume 7, No. 6, June 1984.

Expert-Ease is an expert systems generator which enables people with little or no programming experience to create and interrogate expert models. The concept of expert systems is discussed and Expert-Ease is reviewed.

van Melle, W., "A Domain-Independent System that Aids in Constructing Knowledge-Based Consultation Programs," doctoral thesis, Report No.: STAN-CS-80-820; HPP-80-22, Department of Computer Science, Stanford University, CA, June 1980.

This thesis demonstrates an effective domain-independent system, called EMYCIN, for constructing one class of expert computer programs: rule-based consultants. Such a consultant uses knowledge specific to a problem domain to provide consultative advice to a client. Domain knowledge is represented in EMYCIN primarily as production rules, which are applied by a goal-directed backward-chaining control structure. Rules and consultation data may have associated measures of certainty, and incomplete data are allowed. The system includes an explanation facility that can display the line of reasoning followed by the consultation program or answer questions from the client about the contents of its knowledge base. Other built-in human-engineering features allow the system architect to produce, with a minimum of effort, a consultation program that is pleasing in appearance to the client. To aid the system designer in producing a knowledge base for a domain quickly and accurately, EMYCIN provides a terse, stylized, but easily understood language for writing rules; performs extensive checks to catch common user errors, such as misspellings; and handles all necessary book-keeping chores. To improve efficiency in a running consultation program, EMYCIN provides a rule compiler that transforms the system's production rules into a decision tree, eliminating the redundant computation inherent in a rule interpreter. It then compiles the resulting tree into machine code. The program can thereby use an efficient deductive mechanism for running the actual consultation, while the flexible rule format remains available for acquisition, explanation, and debugging.

Waltz, D. L., "New Tools for Emergency Management," Proceedings of the First Symposium on the Theory and Application of Expert Systems in Emergency Management Operations, Special Publication 717, Gaithersburg, MD: National Bureau of Standards, 1986.

Waterman, D. A., "Exemplary Programming in RITA," in D. Waterman and F. Hayes-Roth (Eds.), Pattern-Directed Inference Systems, New York: Academic Press, 1978.

Waterman, D. A., and F. Hayes-Roth, "An Investigation of Tools for Building Expert Systems," in F. Hayes-Roth, D. A. Waterman, and D. B. Lenat (Eds.), Building Expert Systems, Reading, MA: Addison-Wesley Publishing Company, 1983.

This report describes an investigation into the comparative merits of eight very high-level programming languages designed for building expert systems (i.e., programs that embody the expertise of specialists or scholars in some problem area): EMYCIN, KAS, EXPERT, OPS5, ROSIE, RLL, HEARSAY-III, and AGE. In the investigation, all eight languages were applied to the same environmental crisis-management problem which involved the creation of an on-line assistant to aid in locating and containing an oil or hazardous chemical spill at the Oak Ridge National Laboratory (ORNL). Such an assistant must be able to characterize the spill and its hazards, warn passersby, contain the spill material, locate the spill source to stop the leak, and notify the appropriate authorities. Since the ORNL complex contains over 200 buildings connected by a complicated drainage system that collects all spills, merely locating the spill source is in itself a formidable task.

Waterman, D. A., et al., "Design of a Rule-Oriented System for Implementing Expertise," Report No.: RAND/N-1158-1-ARPA, RAND Corp., Santa Monica, CA, May 1979.

This note describes the preliminary design of a Rule-Oriented System for Implementing Expertise (ROSIE). This system will serve as a tool for model builders seeking to apply expert knowledge to the analysis of problems and the evaluation of solutions in complex domains, especially domains for which useful analytic models are unavailable. Models built within the ROSIE system constitute simulations in the application domain. Control in these models is data-directed. That is, actions are specified by sets of rules having the basic form: 'If CONDITIONS then ACTIONS.' The actions are executed when data matching the conditions appear in the model's data base. Major features of ROSIE include: hierarchic structures on data elements and rules, user selection of rule-invocation strategies, user control of rule iterations, user control of rule and data activation, and sophisticated user support tools.

Weiss, S. M., and C. Kulikowski, "EXPERT: A System for Developing Consultation Models," Proceedings of the Sixth IJCAI, American Association for Artificial Intelligence, 1979.

EXPERT is a system for designing and building models for consultation. An EXPERT model consists of hypotheses (which can be structured into causal and taxonomic networks); findings or observations; and decision rules for logically relating these components. A relatively simple language for describing models is employed. Logical and probabilistic rules are restricted to particular types which implicitly order the tasks performed by the control strategies: classifications, question selection, and explanations. Explicit representations of decision rules are emphasized, as opposed to suboptimal scoring functions. This results in more easily predictable and correctable performance for a model. The system is currently being used to develop consultation models in domains such as rheumatology, ophthalmology, and endocrinology.

Weiss, S. M., and C. Kulikowski, A Practical Guide to Designing Expert Systems, Totowa, NJ: Rowman & Allenheld, 1984.

Wescourt, K. T., and P. W. Thorndyke, "Alternative Knowledge Acquisition Interface Structures," Perceptronics, Inc., Knowledge Systems Branch, Menlo Park, CA, Dec. 1983.

This research developed a design concept for an interactive system to acquire domain knowledge from a training expert. Such a system would facilitate the development of knowledge-based instructional systems by directly eliciting and encoding from domain experts knowledge needed to deliver instruction. An analysis of the process by which knowledge-based systems are constructed indicates that the generality of a knowledge acquisition system must be limited by domain characteristics and by the architecture of the system it serves, and the non-sequential, interacting activities during system development constrain the potential role of automated knowledge acquisition aids. A feasible concept for knowledge acquisition technology, building on prior research in artificial intelligence, involves the notion of class-generic systems for a related set of domains with a fixed architecture and training capabilities. This concept is developed and discussed in the context of the proposed Navy training systems for acquiring models of trainee performance during learning, rules of behavior for an automated opponent in a tactics trainer, and a knowledge base of facts to be subsequently presented to trainees for memorization.

Williams, C., "Software Tool Packages the Expertise Needed to Build Expert Systems," Electronic Design, Volume 32, No. 16, Aug. 1984.

It is claimed that the Automated Reasoning Tool, or ART, is the first comprehensive software package to turn engineers into builders of full-scale expert systems. As such, it contains an inference engine, which resembles the operating system of conventional systems; a knowledge base, which replaces a data base; an editor, which helps create the knowledge base; and a monitor, which lets users visually follow the tool's reasoning steps. Unlike other commercial expert system development packages, ART uses more than one reasoning mechanism. Its language can handle both transformational and logic programming and can pursue both forward and backward chaining. In addition, its knowledge base can be organized by schema and propositions, and through viewpoint mechanisms, the inference engine can consider several hypothetical solutions to a single problem.

Wong, W. G., "LISP for CP/M," Microsystems, August 1983.

Wong, W. G., "PROLOG: The New AI Language Being Used on 5th Generation Japanese Computers Can Now be Run Under CP/M," Microsystems, Jan. 1984.

PROLOG stands for PROgramming in LOGic. It is a logical programming language that has roots in mathematics and logic. Like LISP, it was first used for artificial intelligence research but has since found wider acceptance. In fact, it is being used as the programming language for the Japanese fifth-generation computer project as well as numerous domestic and European research projects. This article presents some of the ideas behind PROLOG, as well as the facilities and performance provided by micro-PROLOG. Prolog is a programming language that differs tremendously from conventional languages such as PASCAL and BASIC because PROLOG is a descriptive language. You tell PROLOG what has to be done, not how it is to be done. This article presents some of the features of PROLOG, as well as a review of a CP/M-based PROLOG--called micro-PROLOG--available from Logic Programming Associates, Ltd.

## 6. USER INTERFACE CONSIDERATIONS

One of the distinct features of an expert system is its ability to explain its actions and decisions to the user. This is a very desirable feature in an expert system, because, unlike a human expert's advice, machine generated advice is seldom accepted without doubt and/or questioning; the system's capacity to explain can help a user to discover when a system is being pushed beyond the bounds of its domain of expertise; the interaction between the user and an expert system is far from what might be called a dialogue, as it is between the user and a human expert; and, since knowledge bases are built incrementally, one is better able to debug the expert system by monitoring what the system is doing.

Today's expert systems accomplish the task of explaining their actions and lines of reasoning through one of the following techniques: (1) By tracing backward from the conclusion in question to the rules producing it. These rules (triggering the conclusion) are then reproduced for the user as an explanation for the action(s) taken or the conclusion(s) drawn. (2) By using previously prepared text (also called canned explanations), i.e., by anticipating the questions and storing the answers as a natural language text. (3) By using the computer code and traces of its execution, and producing explanations directly from it. Each of the above techniques is deficient in one way or another. The first method of reproducing the rules that generated the conclusion can hardly be satisfactory to every user or in every situation. With canned explanations, one would have to know in advance all the questions and answers that might be asked by users of the expert system. In the case of explanations produced from the program code, the quality of the explanations is largely dependent on how the code is written.

### 6.1 DESIRED EXPLANATORY CAPABILITIES

Scott et al. [1977] list the following major goals that must be considered in developing an explanation capability:

(a) The explanation capability must be able to answer all the questions about its knowledge and actions. These include,

- how it made a certain decision
- how it used a piece of information
- what decision it made about some subproblem
- why it did not use a certain piece of information
- why it failed to make a certain decision
- why it required a certain piece of information
- why it did not require a certain piece of information
- how it will find out a certain piece of information (while the consultation is in progress)
- what the system is currently doing (while the consultation is in progress)

(b) The user must be able to get an explanation that answers the question completely and comprehensively.

(c) The explanation capability must be easy to use.

In addition to the above, consideration must also be given to the question of how detailed the explanation should be, and to tailoring the explanation to suit the needs of different users. There are two other areas of development that promise to make expert systems more user friendly: natural language interfaces and intelligent front ends.

## 6.2 TWO APPROACHES FOR MAKING INTERFACES MORE USER FRIENDLY

Use of natural language interface (NLI) is a step toward improving the communication between the user and the expert system. These interfaces allow a user to build and run an expert system by making use of a subset of a natural language like English. Obviously, the single most attractive aspect of having a NLI is the ease of learning how to interact with the system. There are, however, many factors that must be understood before implementing a NLI. First, NLI's are generally more costly to implement because natural languages are less concise compared to the symbolic languages commonly used in building expert systems. Second, they are more expensive to run because of their longer execution times. Third, it is necessary to overcome the ambiguity caused by the imprecise nature of sentences in natural languages. Fourth, the type of users expected to interact with the expert system must be considered; the sophisticated users may be unwilling to compromise with capabilities of an expert system in order to be able to interact with a natural language. Finally, the NLI's can raise the expectations of users regarding their capabilities by leading them to believe that these interfaces are more than just conveniences to the users. For a more detailed discussion on natural language interfaces see Rich [1984].

An intelligent front end (IFE) can be used to build models of user problems by allowing the user to express their problems in familiar languages. The system then translates the resultant problem statement into a language suitable for the software package. The successful implementation of IFE's requires development of techniques that will (a) aid in the knowledge representation (i.e., communication between the user and the software package); (b) solve the problem of synthesizing the software package instructions from the task specification (given by the user); (c) translate the results produced by the software package into the user's language; and (d) carry on a natural language dialogue with the user. In fact, the IFE's are expert systems by themselves, requiring various artificial intelligence tools for building them and implementing them, see Bundy [1984] for further discussion.

## REFERENCES

Anderson, J. R., B. C. Franklin, and B. J. Reiser, "Intelligent Tutoring Systems," Science, Vol. 228, Apr. 1985.

Cognitive psychology, artificial intelligence, and computer technology have advanced to the point where it is feasible to build computer systems that are as effective as intelligent human tutors. Computer tutors based on a set of pedagogical principles derived from the ACT (Advanced Computer Tutoring) theory of cognition have been developed for teaching students to do proofs in geometry and to write computer programs in the language LISP.

Brown, J. S., and R. R. Burton, "A Paradigmatic Example of An Artificially Intelligent Instructional System," International Journal of Man-Machine Studies, Vol. 10, 1978.

This paper describes research directed at understanding and designing artifically intelligent instructional systems that use their knowledge bases and problem-solving expertise to aid the student in several ways. A system is described that derives from the world of manipulatory mathematics based on attribute blocks and examples are given of interaction with it. The system can answer student questions and evaluate his theories as well as critize his solution paths. It can also form structural models of his reasoning strategies to identify his fundamental misconceptions and determine when and how to provide remediation, hints or further instruction.

Bundy, A., "Intelligent Front Ends," in M. A. Bramer (Ed.), Research and Development in Expert Systems: Proceedings of the Fourth Technical Conference of the British Computer Society Specialist Group on Expert Systems, Cambridge, England: Cambridge University Press, 1984.

Carroll, J. M., "How to Turn an Emergency Simulation Into a "Video Game" for Training," Proceedings of the Conference on Computer Simulation in Emergency Planning, La Jolla, CA: Simulation Councils, Inc., Jan. 1983.

Computer simulation can be used to train decision makers. The programs should be interactive and are readily accepted by trainees if they resemble video games. Sometimes the structure of the problem already resides in a simulation program used for planning and resource allocation. This paper tells how to convert a simulation program for allocating fire-fighting resources into a microcomputer-based training exercise for fire dispatchers.

Clancy, W. J., "Extensions to Rules for Explanation and Tutoring," in B. G. Buchanan and E. H. Shortliffe (Eds.), Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project, Reading, MA: Addison-Wesley Publishing Company, 1984.

Clancey, W. J., "Tutoring Rules for Guiding a Case Method Dialogue," International Journal of Man-Machine Studies, Vol. 12, 1980.

The first version of an intelligent computer-aided instruction program built on MYCIN-like expert systems has been implemented. This program, named GUIDON, is a case method tutor in which the problem-solving and tutorial dialogue capabilities are distinct. The expertise to be taught is provided by a rule-based consultation program. The dialogue capabilities constitute teaching expertise for helping a student solve a case. This paper describes the rule-based formalism used by MYCIN-like programs and then argues that these programs are not sufficient in themselves as teaching tools. We have chosen to develop a mixed-initiative tutor that plays an active role in choosing knowledge to present to a student based on his competence and interests. It is argued that is desirable to augment the domain expertise of MYCIN-like programs with other levels of domain knowledge that help explain and organize the domain rules. Finally, the author claims that it is desirable to represent teaching expertise explicitly, using a flexible framework that makes it possible to easily modify tutorial strategies and communicate them to other researchers. The design of the GUIDON program is based on natural language studies of discourse in AI. In particular, the framework used integrates domain expertise in tutorial dialogues via explicit, modular tutoring rules that are controlled by a communication model. This model is based on consideration of the student's knowledge and interests, as well as the tutor's plans for the case session. This paper discusses interesting examples of tutoring rules for guiding discussion of a topic and responding to a student's hypothesis based on the evidence he has collected.

Cullingford, R. E., et al., "Automated Explanations as a Component of a Computer-Aided Design System," IEEE Transactions on Systems, Man, and Cybernetics, SMC-12(2), 1982.

Davis, R., "Interactive Transfer of Expertise: Acquisition of New Inference Rules," Artificial Intelligence, Vol. 12, 1979.

Eden, C., "A System to Help People Think About Problems," Computer Weekly, No. 876, Sept. 1983.

For the past five years, a group of academics in the school of management at the University of Bath has been developing software for combining expertise. Their focus has been on using computer software to help a team of decision makers work on complicated problems -- the sort of problems that demand good use of the expertise of each member of the team, and where that expertise is mostly based on experience rather than expert knowledge. The author describes how users can argue with an expert system by adding conflicting points of view.

Fitter, M., "Towards More 'Natural' Interactive Systems," *International Journal of Man-Machine Studies*, Vol. 12, 1980.

It is argued that to obtain the maximum benefit from interactive computer systems, principles of program and dialogue design are needed. It is unlikely that natural languages such as English will provide a suitable basis for designing a man-computer dialogue. The objective should be to model the task domain in a way that will be comprehensible to the user and to provide an explicit image of the underlying processes. Design principles are discussed both for general purpose programming languages and bespoke languages intended as tools for a specific purpose. It is concluded that AI research may eventually provide intelligent tools with the inferential powers necessary for a genuine dialogue, but that for the time being it is better to make the underlying mechanisms and their limitations as explicit as possible.

Gable, A., and C. V. Page, "The Use of Artificial Intelligence Techniques in Computer-Assisted Instruction: An Overview," *International Journal of Man-Machine Studies*, Vol. 12, 1980.

One of the major goals of research in artificial intelligence is the representation of knowledge so that a computer can solve problems or communicate in a manner which exhibits common sense. Few programs for computers, including those for education, possess behavior which approaches any facet of the constellation of human skills and knowledge which are imprecisely called common sense. However, the revolutionary decline in hardware costs now makes it possible to consider economically viable, sophisticated designs for computer-aided instruction systems possessing some of the common sense attributes of a human tutor. In this survey, we examine, in depth, techniques from aritifical intelligence that can be used to endow a Computer-Aided Instruction system with approximations to some of the desirable qualitities of a human tutor. We consider both techniques which have been proved in prototype systems for Computer-Aided Instruction and some techniques which were originally developed for other purposes.

Gerring, P., E. H. Shortliffe, and W. Van Melle, "The Interviewer/Reasoner Model: An Approach to Improving System Responsiveness in Interactive AI Systems," *AI Magazine*, Vol. 3, No. 4, 1982.

Grishman, R., "Natural Language Processing," *Journal of the American Society for Information Science*, Vol. 35, No. 5, Sept. 1984.

Kidd, A. L., "Human Factors in Expert Systems," *Proceedings of the Ergonomics Society's Conference*, London, England: Taylor and Francis, Mar. 1983.

The development of expert or intelligent knowledge-based systems is raising some new and very important human factor issues in the field of man-computer interaction. With the advent of thinking, reasoning systems, the cognitive characteristics of the human mind must become a critical design constraint in the production of such systems. This paper attempts to identify the major research questions in this area.

Mark, W., "Natural Language Help in the Consul System," Proceedings of the 1982 AFIPS National Computer Conference, Vol. 51, Arlington, VA: AFIPS Press, June 1982.

If one uses the model of asking an expert, it is fairly clear what users want from a help system: a service that tells them how to do something they want to do. But current help systems are not like this. They can tell the users about system capabilities, but not in relation to what they want to do. Most help systems are simply databases of on-line documentation-system manuals, not system experts. Like system manuals, using them to figure out how to do something or to figure out what went wrong is a last resort--when no one else is around. Providing real expert help requires reasoning in terms of models of what the user wants to do and what the system can do. These models also make it possible to provide facilities for natural language understanding and generation. Thus, users can deal with the system in much the same way that they deal with a human expert: by asking questions and receiving advice in English.

Mikulich, L. I., "Natural Language Dialog Systems: A Pragmatic Approach," in J. E. Hayes, D. Michie and Y. H. Pao (Eds.), Machine Intelligence 10, Chichester, England: Norwood, 1982.

An approach is described for the development of natural language interfaces with data bases. Pragmatic-oriented analysis of utterences made by casual users is a key concept of the approach. An example of applying is considered using the DISPUT family of systems.

Miller, P. L., et al., "Teaching With 'Attending': A Practical Way to 'Debug' an Expert Knowledge Base," AAMSI Congress 83 Proceedings, Bethesda, MD: American Association of Medical Systems and Informatics, May 1983.

O'Shea, T., "Rule-Based Computer Tutors," in D. Michie (Ed.) Expert Systems in the Micro-Electronic Age, Edinburgh, Scotland: Edinburgh University Press, 1979.

Two important components of any computer tutor are the tutorial strategy and the student model. Production Rules (PRs) can usefully be employed in the design and implementation of both these components. The use of PRs forces careful analysis of the taxonomic aspects of computer tutor design. PRs are modular, explicit, amenable to automatic manipulation and hence, applicable to the construction of adaptive tutors. Their use provides a valuable link to work on the analysis of human information processing strategies.

Pollack, M. E., J. Hirschberg, and B. Weber, "User Participation in the Reasoning Process of Expert Systems," Proceedings of the National Conference on Artificial Intelligence, AAAI-82, American Association for Artificial Intelligence, 1982.

The authors argue that expert sytems, if they are to satisfy the legitimate needs of their users, must include dialogue capabilities as sophisticated as those proposed in current natural language research. In particular, they must allow the user to direct the flow of the dialogue and must take into account the user's goals and expectations both in analyzing the user's statements and in providing appropriate responses. Their claims are corroborated by an analysis of transcripts of a naturally occurring expert system, a radio talk show in which callers ask an expert for financial advice. The authors present data demonstrating that user-expert dialogues are best viewed as a negotiation process, and describe the exchanges that compose the dialogue in terms of the motivations, goals, strategies, and moves of the participants.

Reboh, R., "Using a Matcher to Make an Expert Consultation System Behave Intelligently," Proceedings of the National Conference on Artificial Intelligence, AAAI-80, American Association for Artificial Intelligence, 1980.

Rich, E., "Natural Language Interfaces," Computer, Vol. 17, No. 9, Sept. 1984.

Richie, G. N., "ATLANTIS: A Disaster Simulated Exercise for Training Purposes," Proceedings of the Conference on Computer Simulation in Emergency Planning, La Jolla, California: The Society for Computer Simulation, Jan. 1983.

This paper examines the process of identifying training needs and methodology, and the selection and development of a computer-based simulation exercise as a means of meeting these requirements. In doing so, it seeks to identify guidelines which may be helpful to those involved in similar projects. These guidelines relate to conception, identification, and establishment of modeling parameters, project management and development trials, and the integration of the simulation exercise into a study program. The paper is based on the experience of the author in a research project which examined the role and responsibilities of the administration in maintaining life support systems in disaster situations in developing countries.

Schachter-Radig, M. J., "Prospects for the Use of Knowledge-Based Systems - Demands on the User," Proceedings of the COMPAS '84, Berlin, Germany: VDE-Verlag, Oct. 1984.

Scott, A. C., et al., "Explanation Capabilities of Knowledge-Based Production Systems," American Journal of Computational Linquistics, Microfiche 62, 1977.

Shrager, J., and T. Finin, "An Expert System That Volunteers Advice," Proceedings of the National Conference on Artificial Intelligence, AAAI-82, American Association for Artificial Intelligence, 1982.

The system described in this paper essentially encodes the pattern recognition knowledge of a consultant and applies these patterns to user input. When inefficiencies are recognized, the automatic observer can point the way to a more effective use of system capabilities by referring the user to the help system or some other expert. Otherwise, it can generate a help message of its own using the context of the particular sequence to construct useful dialogue. This work represents an attempt to extend the range of interactive advisor systems. User aids currently in service help only users who cause errors which would invoke an error recovery system or those who know how and when to ask for help. The authors have provided a means by which the system can automatically tutor intermediate level users who do not make trivial errors but who are not using commands in an effective way. Therefore, it can bridge the gap between the introductory user aids and more technically-oriented expert advisors.

Sleeman, D and J. S. Brown (Eds.), Intelligent Tutoring Systems, London, England: Academic Press, 1982.

Smith, R. G., et al., "Declarative Task Description as a User-Interface Structuring Mechanism," Computer, Vol. 17, No. 9, Sept. 1984.

Swartout, W. R., "Explaining and Justifying Expert Consulting Programs," in W. J. Clancey and E. H. Shortliffe (Eds.), Readings in Medical Artificial Intelligence: The First Decade, Reading, MA: Addison-Wesley Publishing Company, 1984.

Ueno, H., "An End-User-Oriented Language to Develop Knowledge-Based Expert Systems," Proceedings of the COMPCON 83 Fall, Silver Spring, MD: Computer Society Press, Oct. 1983.

A description is given of the COMEX (Compact Knowledge-Based Expert System) expert system language for application-domain users who want to develop a knowledge-based expert system by themselves. The COMEX is being used in several application domains such as medicine, education and industry.

Wallis, J. W., and E. H. Shortliffe, "Customized Explanations Using Causal Knowledge," in B. G. Buchanan and E. H. Shortliffe (Eds.), Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project, Reading , MA: Addison-Wesley Publishing Company, 1984.

Waterman, D. A., "User-Oriented Systems for Capturing Expertise: A Rule-Based Approach," in D. Michie (Ed.), Expert Systems in the Micro-Electronic Age, Edinburgh, Scotland: Edinburgh University Press, 1979.

User-oriented systems for capturing expertise are discussed and illustrated within the context of RITA and ROSIE, two rule-based systems developed for building intelligent interfaces and modeling expert knowledge.

Woolf, B., and D. D. McDonald, "Building a Computer Tutor: Design Issues," Computer, Vol. 17, No. 9, Sept. 1984.

In this article, the authors discuss how a deep understanding of a student can be constructed in an artificial intelligence program and how this understanding, coupled with a facility for language generation, can be used to build a flexible machine tutor. An effective tutor must deal with a fundamental problem of communication: to determine how messages are received and understood, and to formulate appropriate answers. A human tutor, therefore, takes the time to double-check and reviews a student's knowledge to find out whether or not he understands what the tutor has said. Compounding the situation is the fact that a student is generally unaware of what he does not know. This means that a tutor, more than a typical speaker, must verify that both parties know what information has been covered, what is missing, and which communication might be erroneous.

# 7. DEALING WITH UNCERTAINTY IN THE DESIGN, DEVELOPMENT AND USE OF EXPERT SYSTEMS

## 7.1 IMPORTANCE OF DEALING WITH UNCERTAINTY IN EXPERT SYSTEMS

Expert systems need to be designed to work with data and information that is uncertain, incompatible, and incomplete. There are several reasons for this. First, many types of information, some of it inexact, cannot be encoded precisely in a machine usable form. Second, because of its subjective nature, the expertise captured from human experts cannot always be faithfully translated into actionable machine code. Third, users of expert systems are often unable to give all the information and data required by the expert system. Fourth, for many decisions, a combination of uncertain data and information is common. Finally, the value of an expert system would be considerably diminished if it always required complete and certain data.

## 7.2 THE HANDLING OF UNCERTAINTY IN EXPERT SYSTEMS

Expert systems must be able to make acceptable decisions even when there is uncertainty regarding the data and information, or even the expertise. Human experts are capable of using a number of different techniques to deal with uncertainty. In contrast, an expert system is tied to a single method for handling uncertainty. There are four methods of dealing with uncertainty in expert systems that have received varying degrees of acceptance. These methods are based on the following theories: Bayesian probability theory, certainty theory, possibility theory and fuzzy logic, and the Dempster-Shafer theory of evidence. Among these, Bayesian probability theory and certainty theory are the only ones that have been widely used. Although much information is available on the theory and applicability of fuzzy logic, its usefulness has not been fully exploited. The Dempster-Shafer theory has only recently received attention.

Each of these methods differs in the type and amount of data or information needed to use it. For example, assuming that the knowledge in an expert system is represented by If-Then rules (as it is in rule-based expert systems), the Bayesian model would require values for (a) prior probabilities (i.e., probabilities without any evidence either supporting or negating some conclusion), and (b) conditional probabilities (i.e., probabilities associated with conclusions being true, given some piece of information or evidence) for all conclusions that can be considered by the system. If the same expert system were to use the certainty theory, the worth of each piece of evidence would be expressed as a "measure of increased belief" or a "measure of increased disbelief" (values between 0 and 1) in a conclusion, depending on whether that piece of evidence increases or decreases the truth of a conclusion. In both instances, the measures would be subjectively assigned by an expert in the problem domain, see Duda [1983] (Chapter 3) for further discussion on types of data used by different models in dealing with uncertainty. We next give a brief discussion of each of these four methods.

## Bayesian Probability Theory

This theory enables the expert system to express its confidence in any conclusion or decision made as a conditional probability statement of the following form:

> Given the specified evidence, C can be concluded with a probability of 0.8.

If the conclusion C had been certain (based on the specified evidence), a probability value of 1.0 would have been attached to it. Similarly, if the evidence had led to concluding 'not C' with certainty, a probability value of zero would have been attached to the above statement. All other probability values (falling between 0.0 and 1.0) attached to such conditional statements give the degrees of uncertainty associated with them. Essentially, Bayes' Theorem can be seen as a mechanism to calculate the probability of a conclusion, in light of specified evidence, from the a priori probability of the conclusion and the conditional probabilities relating the observations to the conclusions in which they are expected. Application of the Bayesian theory involves the following:

(a) Each If-Then rule in the knowledge base is represented in the form of a conditional probability statement (i.e., a probability value is attached to every conclusion given a possible evidence).

(b) It is assumed that the portion of evidence (represented as a probability value) that does not support a conclusion, supports its negation.

(c) If no information is available regarding a needed piece of evidence, the probability of a conclusion being true is equal to its prior probability.

(d) The general Bayesian probability rule for computing conditional probability is applied to combine the various pieces of evidence and for assigning the conditional probability to the conclusion drawn.

(e) A check for agreement between conditional probabilities derived using the Bayesian rule and those obtained from the expert, if he were to independently assess the truth of a conclusion given the various pieces of evidence, is carried out.

Although computationally simple, the usefulness of Bayesian approach is limited by practical difficulties, principally by the lack of data adequate to estimate accurately the prior and conditional probabilities. The conclusions under consideration are often assumed to be mutually exclusive and exhaustive (i.e., only one of the several possible conclusions is entertained, given some set of evidence). The set of observations is also assumed to have independent elements, although in practice, Bayesian Theory does not require this. The model does not generally allow for changes in observations (and hence conclusions) over time. Revisions of prior and conditional probabilities to achieve consistency between estimates of the expert and those calculated using Bayesian rule can be quite time consuming. Finally, problems also arise when two or more pieces of evidence independently point to the same conclusion, but together rule out the conclusion. An example of this could be the situation where E1 suggests C with weight W1 and E2 suggests C with weight W2, but the presence of both E1 and E2 rules out C. One way in which such cases are handled is by using an additional rule of the form E1 & E2 suggest C with weight W. See Adams [1984] and Duda, Hart and Nilsson [1976] for more detailed discussion on how to use Bayesian probability theory in dealing with uncertainty.

## Certainty Theory

Certainty theory, developed and first put to use by Shortliffe and Buchanan [1975] in their MYCIN medical consultation system, is currently the most widely applied methodology in dealing with uncertainty in expert systems. Central to this theory is the concept of a certainty factor (CF), a number between -1 and 1 attached to every rule, representing the degree of belief that the rule implies for its assertion (or the conclusion). The certainty factor itself is the difference between "the measure of increased belief in conclusion C, given evidence E" (MB), and "the measure of increased disbelief in conclusion C, given the same evidence E" (MD). Both MB and MD are numbers between 0 and 1. The definition of certainty factor brings out the important difference between this approach and the Bayesian approach. The Bayesian approach specifies that if we know the value of evidence E in favor of some conclusion C to be x (i.e., conditional probability of a conclusion C being true given evidence E), then the value of the evidence E that does not support conclusion C is given by 1-x (i.e., conditional probability of a conclusion not being true, given evidence E). The certainty theory uses the notion that one piece of evidence cannot both favor and disfavor a single assertion. Thus, if for any piece of evidence the value of MB is greater than 0, the value of MD is taken to equal 0, and vice versa. A typical rule in a system using certainty theory is represented as follows (taken from Shortliffe and Buchanan [1975]):

IF: 1) the stain of the organism is gram negative, and
2) the morphology of the organism is rod, and
3) (a) the aerobicity of the organism is aerobic, or
(b) the aerobicity of the organism is unknown,

THEN: there is suggestive evidence (.7) that the class of the organism is enterobacteriaceae

Application of certainty theory in expert systems consists of revising or updating uncertainties as the rules are applied and new evidence becomes available. Initially, the certainty associated with any assertion is 0. If a rule says E1 implies C with a certainty factor x, then the certainty of C is changed to x when evidence E1 is observed. Finally, as new pieces of evidence become available, they are combined using the procedures specified by the theory.

At this point, some important observations regarding this theory are in order. First, the amount of data required to use this theory is of the same order as in case of the Bayesian approach. Second, although much of the reasoning in this theory parallels that of the probability theory, the probability rules cannot be strictly applied to its concepts. Finally, there are instances where applications of this theory yield counterintuitive results because of the way some of its concepts are defined, see Adams [1984] for more details. More recently, Shortliffe and Buchanan [1984] report changes that have been introduced in the definitions to correct this deficiency. A careful delineation of this theory's relationship to conventional probability theory can contribute to a better understanding of its assumptions and approximations.

### Possibility Theory and Fuzzy Logic

The main premise upon which this theory is built may be stated as follows: The random model of the probability theory is more suited for problems involving the measure of information, but it is inappropriate for problems where the meaning of information must be discerned. The problem of dealing with imprecision is very important in the area of expert systems, as one rarely has the luxury of being able to work with exact definitions, descriptions, or assertions. In possibility theory, an attempt is made to formalize vagueness and imprecision, just as probability theory is used as a formalism in treating randomness. The theory of fuzzy sets (upon which the possibility theory is based) allows us to express all kinds of imprecise entities by specifying to what extent they have membership in suitably labeled fuzzy sets. For example, if one does not know exactly the probability of an event, (e.g., it will rain tomorrow) but is able to say that it is either very probable or not very probable, then we can establish corresponding fuzzy subsets. We can then decide to what degree we are willing to admit the individual elements representing evidence into these sets. In other words, a membership function is established which associates with each element a number in the interval [0,1] which represents the degree of its membership in the fuzzy set in question. By introducing the concept of degree of membership, this theory removes the rigidity of ordinary set theory of either admitting or rejecting an element (with no possibility of it partially belonging to the set).

The usefulness of fuzzy set theory in dealing with uncertainty in expert systems arises out of its being able to deal with uncertainty due to vagueness, where probability theory is inappropriate. Although some efforts have been made in applying this theory to expert systems, much of it is in the developmental stage. There is, however, ample literature for the interested reader (Baldwin [1979, 1983], Bellman and Giertz [1973], Bellman and Zadeh [1970], Dubois and Prade [1979, 1985], Green [1984], Negoita [1985], Zadeh [1980, 1982, 1983a, 1983d, 1986], and Zimmerman [1984]).

### The Dempster-Shafer Theory of Evidence

The foundations of this theory were laid over two decades ago by Dempster [1963, 1967, 1968] and extended by Shafer [1975, 1976]. It is only recently that researchers have shown a renewed interest in this theory and have investigated its applicability as a model for dealing with uncertainty in expert systems, see Barnett [1981], Dillard [1983], Friedman [1981], Garvey, Lowrence and Fischler [1981], Gordon and Shortliffe [1984].

The main advantage of the Dempster-Shafer theory is its ability to allow the system to narrow down the number of hypotheses that must be considered in decision making. This property is especially valued by builders of diagnostic expert systems in areas like medicine. In a Bayesian probabilistic model, if some evidence or data is found that only partially supports a certain hypothesis being tested, the portion not supporting it is taken as evidence against the hypothesis being true, which may not be necessarily true. The Dempster-Shafer theory is not weakened by such an assumption. Further, unlike the Bayesian approach which does not distinguish between uncertainty, or lack of knowledge, and equal certainty, this theory more faithfully reflects the evidence-gathering process by using such evidence to narrow down the original hypothesis set to a different subset.

In closing, it is useful to view each of these approaches as proposing a tool for obtaining the desired behavior in an expert system, and not how the information must be processed.

## REFERENCES

Adams, J. B., "Probabilistic Reasoning and Certainty Factors," in B. G. Buchanan and E. H. Shortliffe, (Eds.) Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project, Reading, MA: Addison-Wesley Publishing Company, 1984.

Baas, M. S., and H. Kwakernaak, "Rating and Ranking of Multiple-Aspect Alternatives Using Fuzzy Sets," Automatica, Volume 13, 1977.

Baldwin, J. F., "Fuzzy logic and Its Application to Fuzzy Reasoning," in M. M. Gupta, R. K. Ragade, and R. R. Yager (Eds.), Advances in Fuzzy Set Theory and Applications, Amsterdam, Netherlands: North-Holland Publishing Company, 1979.

Baldwin, J. F., "A Fuzzy Relational Inference Language for Expert Systems," Proceedings of the Thirteenth International Symposium on Multiple-Valued Logic, New York: IEEE, May 1983.

A fuzzy relational inference language, FRIL, is discussed in relation to its application for designing intelligent knowledge bases and expert systems. It is a high level query language with automatic reasoning and is based upon fuzzy set and relation theories.

Barnett, J. A., "Computational Methods for a Mathematical Theory of Evidence," Proceedings of the Seventh IJCAI, American Association for Artificial Intelligence, 1981.

Many knowledge-based expert systems employ numerical schemes to represent evidence, rate competing hypotheses, and guide search through the domain's problem space. This paper has two objectives: first, to introduce one such scheme, developed by Dempster and Shafer, to a wider audience; second, to present results that can reduce the computation-time complexity from exponential to linear, allowing this scheme to be implemented in many more systems. In order to enjoy this reduction, some assumptions about the structure of the type of evidence represented and combined must be made. The assumption made here is that each piece of the evidence either confirms or denies a single proposition rather than a disjunction. For any domain in which the assumption is justified, the savings are available.

Bellman, R. E. and L. A. Zadeh, "Decision-Making in a Fuzzy Environment," Management Science, Vol. 17, No. 4, Dec. 1970.

By decision-making in a fuzzy environment is meant a decision process in which the goals and/or constraints, but not necessarily the system under control, are fuzzy in nature. This means that the goals and/or constraints constitute classes of alternatives whose boundaries are not sharply defined. Fuzzy goals and fuzzy constraints can be defined precisely as fuzzy sets in the space of alternatives. A fuzzy decision, then, may be viewed as an intersection of the given goals and constraints. A maximizing decision is defined as a point in the space of alternatives at which the membership function of a fuzzy decision attains its maximum value. The use of these concepts is illustrated by examples involving multi-stage decision processes in which the system under control is either deterministic or stochastic. By using dynamic programming, the determination of a maximizing decision is reduced to the solution of a system of functional equations. A reverse-flow technique is described for the solution of a functional equation arising in connection with a decision process in which the termination time is defined implicitly by the conditions that the process stops when the system under control enters a specified set of states in its state space.

Bellman, R. E., and M. Giertz, "On the Analytic Formalism of the Theory of Fuzzy Sets," Information Sciences, Vol. 5, 1973.

Bezdek, J. C., "Numerical Taxonomy with Fuzzy Sets," Journal of Mathematical Biology, Vol. 1, No. 1, 1974.

Bezdek, J. C., Pattern Recognition with Fuzzy Objective Function Algorithms, New York: Plenum Press, 1981.

Carlson, C., "Tackling an MCDM-Problem with the Help of Some Results from Fuzzy Set Theory," European Journal of Operational Research, Vol. 10, 1982.

The problem to be addressed and tackled in this paper arose as a byproduct from some efforts at solving problems involving multiple goals by linking linear and goal programming models. The critical issue was that some forms for interdependence among the goals could not be handled in the programming models. Here the author deals with a set of goals -- with realistic counterparts in a Finnish plywood industry -- in which a subset of the goals are (i) conflicting, (ii) unilaterally supporting and (iii) mutually supporting. It is furthermore observed that the elements of studied set of goals may be partly independent and partly inter-dependent, which makes the context a fullfledged MCDM-problem. It is tackled with a technique which is based on the theory of fuzzy sets, the conceptual framework for fuzzy decisions and the algorithms developed for fuzzy mathematical programming. The resulting fuzzy multi-objective programming model is simplified and tested with the help of a fairly complex numerical example.

Chang, C. L., "Decision Support in an Imperfect World," Proceedings of Trends and Applications 1983 on Automating Intelligent Behavior: Applications and Frontiers, Silver Spring, MD: IEEE Computer Society Press, May 1983.

By a decision support system, the author means an expert system that the user can use to inquire information to make his decision. Such a system will be based on an expert knowledge base. It is believed that the knowledge base is more than facts and rules. It may include less tangible and less codifiable factors like opinions, judgments, educated guesses, as well as factual information and logic rules for reasoning. That is, the knowledge may be explicit, logical, heuristic, or fuzzy. This paper presents some methods to add fuzzy information into a relational database. Specifically, it considers the treatments of fuzzy queries. An answer to a fuzzy query will be a fuzzy set from which a further detailed study can be made.

Dempster, A. P., "On Direct Probabilities," *Journal of the Royal Statistical Society*, Series B, Vol. 25, 1963.

The concept of a direct probability argument is introduced. This is a means of assigning distributions a posteriori which satisfies the requirement that such distributions shall be valid on the observations, but which differs from the inverse probability argument in that distributions a priori are not needed. Direct probabilities are closely related to fiducial probabilities. This theory requires a generalized concept of sampling a family F of distributions which is discussed in Section 3. Direct inference from a sample of n is discussed in Section 4 for a single continuous parameter and in Section 5 for non-parametric cases.

Dempster, A. P., "Upper and Lower Probabilities Induced by a Multivalued Mapping," *Annals of Mathematical Statistics*, Vol. 38, 1967.

Dempster, A. P., "New Methods for Reasoning Towards Posterior Distributions Based on Sample Data," *Annals of Mathematical Statistics*, Vol. 37, 1967.

This paper redefines the concept of sampling from a population with a given parametric form, and thus, leads up to some proposed alternatives to the existing Bayesian and fiducial arguments for deriving posterior distributions. Section 2 spells out the basic assumptions of the suggested class of sampling models, and Section 3 suggests a mode of inference appropriate to the sampling models adopted. A novel property of these inferences is that they generally assign upper and lower probabilities as given by Bayesian or fiducial arguments. Sections 4 and 5 present details of the new arguments for binomial sampling with a continuous parameter p and for general multinomial sampling with a finite number of contemplated hypotheses. Among the concluding remarks, it is pointed out that the methods of Section 5 include as limiting cases situations with discrete or continous observables and continuously ranging parameters.

Dempster, A. P., "A Generalization of Bayesian Inference," *Journal of the Royal Statistical Society*, Series B, Vol. 30, No. 2, 1968.

Procedures of statistical inference are described which generalize Bayesian inference in specific ways. Probability is used in such a way that in general only bounds may be placed on the probabilities of given events, and probability systems of this kind are suggested both for sample information and for prior information. These systems are then combined using a specified rule. Illustrations are given for inferences about trinomial probabilities, and for inferences about a monotone sequence of binomial $p_i$. Finally, some comments are made on the general class of models which produce upper and lower probabilities and on the specific models which underlie the suggested inference procedures.

Dillard, R. A., "Computing Confidences in Tactical Rule-Based Systems by Using Dempster-Shafer Theory," Report No.: NOSC/TD-649, Naval Ocean Systems Center, San Diego, CA, Sept. 1983.

This report describes the implementation of confidence computing methods based on Demspter-Shafer Theory. The theory is applicable to tactical decision problems that can be formulated in terms of sets of mutually exclusive and exhaustive propositions. Dempster's combining procedure, a generalization of Bayesian inference, is used to combine probability mass assignments supplied by independent bodies of evidence. The computing procedures are implemented in the rule-based system ROSIE, and apply to all valid mass assignments. An ordering strategy is used to combine various kinds of assignments by using different procedures that exploit the special features of each. Applications to platform typing and contact association are demonstrated.

Dubois, D., and H. Prade, "Outline of Fuzzy Set Theory: An Introduction," in M. M. Gupta, R. K. Ragade, and R. R. Yager (Eds.), Advances in Fuzzy Set Theory and Applications, Amsterdam, Netherlands: North-Holland Publishing Company, 1979.

Dubois, D., and H. Prade, "Decision-Making Under Fuzziness," in M. M. Gupta, R. K. Ragade, and R. R. Yager (Eds.), Advances in Fuzzy Set Theory and Applications, Amsterdam, Netherlands: North-Holland Publishing Company, 1979.

Duda, R. O., P. E. Hart and N. J. Nilsson, "Subjective Bayesian Methods for Rule-Based Inference Systems," Technical Note No. 124, Stanford Research Institute, Menlo Park, CA, Jan. 1976.

The general problem of drawing inferences from uncertain or incomplete evidence has invited a variety of technical approaches, some mathematically rigorous and some largely informal and intuitive. Most current inference systems in artificial intelligence have emphasized intuitive methods, because the absence of adequate statistical samples forces a reliance on the subjective judgment of human experts. This paper describes a subjective Bayesian inference method that realizes some of the advantages of both formal and informal approaches. Of particular interest are the modifications needed to deal with the inconsistencies usually found in collections of subjective statements.

Erman, L. D., et al., "The Hearsay-II Speech Understanding System: Integrating Knowledge to Resolve Uncertainty," Computing Surveys, Vol. 12, 1980.

Fordon, W. A., and J. C. Bezdek, "The Application of Fuzzy Set Theory to Medical Diagnosis," in M. M. Gupta, R. K. Ragade, and R. R. Yager (Eds.), Advances in Fuzzy Set Theory and Applications, Amsterdam: North-Holland Publishing Company, 1979.

Friedman, L., "Extended Plausible Inference," Proceedings of the Seventh IJCAI, American Association for Artificial Intelligence, 1981.

Gaines, B. R., "Foundations of Fuzzy Reasoning," International Journal of Man-Machine Studies, Vol. 8, 1976.

This paper gives an overview of the theory of fuzzy sets and fuzzy reasoning as proposed and developed by Zadeh. In particular, it reviews the philosophical and logical antecedents and foundations for this theory and its applications. The extension of basic set operations to such fuzzy sets, and the relationship to other multivalued logics for set theory, are then outlined. The fuzzification of mathematical structures leads naturally to the concepts of fuzzy logic and inference, and consideration of implication suggests Lukasiewicz infinite-valued logic as a base logic for fuzzy reasoning. The paradoxes of the barber, and of sorites, are then analysed to illustrate fuzzy reasoning in action and lead naturally to Zadeh's theory of linguistic hedges and truth. Finally, the logical, model-theoretic and psychological derivations of numeric values in fuzzy reasoning are discussed, and the rationale behind interest in fuzzy reasoning is summarized.

Gaines, B. R., and L. J. Kohout, "The Fuzzy Decade: A Bibliography of Fuzzy Systems and Closely Related Topics," International Journal of Man-Machine Studies, Vol. 9, 1977.

The main part of this paper consists of a bibliography of some 1150 items, each keyword-indexed with some 750 being classified as concerned with fuzzy system theory and its applications. The remaining items are concerned with closely related topics in many-valued logic, linguistics, the philosophy of vagueness, etc. These background references are annotated in an initial section that outlines the relationship of fuzzy system theory to other developments and provides pointers to various possible fruitful interrelationships. Topics covered include: the philosophy and logic of imprecision vagueness, other non-standard logics, foundations of set theory, probability theory, fuzzification of mathematical systems, linguistics and psychology, and applications.

Garvey, T. D., J. D. Lowrence, and M. A. Fischler, "An Inference Technique for Integrating Knowledge from Disparate Sources," Proceedings of the Seventh ICJAI, American Association for Artificial Intelligence, 1981.

Gordon, J., and E. H. Shortliffe, "The Dempster-Shafer Theory of Evidence," in B. G. Buchanan and E. H. Shortliffe, Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project, Reading, MA: Addison-Wesley Publishing Company, 1984.

Gougen, Jr., J. A., "Concept Representation In Natural and Artificial Languages: Axioms, Extensions and Applications for Fuzzy Sets," International Journal of Man-Machine Studies, Vol. 6, 1974.

This paper reports research related to mathematics, philosophy, computer science and linguistics. It gives a system of axioms for a relatively simple form of fuzzy set theory and uses these axioms to consider the accuracy of representing concepts in various ways by fuzzy sets.

Byproducts of this approach include a number of new operations and laws for fuzzy sets, parallel to those for ordinary sets, and a demonstration that all the basic operations are intrinsically determined. In addition, the paper explores both hierarchical and algorithmic extensions of fuzzy sets, and then applications to problems in natural language semantics and combinatorics. Finally, the paper returns to the problem of representing concepts, and discusses some implications for artificial intelligence.

Green, N. L., "Decision Support With Fuzzy Technology," Proceedings of the Sixteenth Annual Meeting of the American Institute for Decision Sciences, Atlanta, GA: American Institute for Decision Sciences, Nov. 1984.

In strategic planning, much significant information is imprecise, yet choosing a strategy is an important decision meriting DSS. Fuzzy sets technology makes computer support possible using strategic manager's natural language descriptions of the firm and its environment. An innovative DSS augments a manager's information processing capability using a fuzzy expert system to search the database for possible strategies. It encourages the decision maker's pattern recognition skills, intuition, and creativity to produce additional strategies inspired by visual displays. The system leaves to the user's judgement the choice of a strategy from the alternatives generated by system and user together.

Humphreys, P., "Knowledge and Probability in Decision Support: Handling Uncertainty in Structuring Ill-defined Problems," IEE Colloquium on Decision Support Aspects of Expert Systems, (Digest No. 67), London, England: IEE, May 1984.

The author examines factors which need to be taken into account in providing decision support for decision makers facing ill-defined decision problems; that is, those problems where at the outset there is considerable uncertainty about what is involved in the problem and how to represent it. It reviews ways in which decision-theoretical models take four different types of uncertainty into account. Within each of these structural variants, three levels of problem representation are identified. But it is shown how at least a further two levels must be taken into account in handling procedural uncertainty when structuring ill-defined problems. Operations at these higher levels are characterized in terms of the development of problem structuring languages and the exploration of small worlds. A multi-level problem structure generating scheme is proposed. Examples of typical properties of real-life decision problems represented at each level are given.

Kohout, L. J., "Fuzzy Decision Making and Its Impact on the Design of Expert Systems," IEE Colloquium on Decision Support Aspects of Expert Systems, (Digest No. 67), London, England: IEE, 1984.

Some methods for decision making concerned with essentially imprecise data are discussed. The discussion is illustrated by the examples from those domains of medical decision making in which classical probabilistic and statistical methods fail. Fuzzy sets and relations that are the essential

tools of this approach have strong impact on inference in expert systems, as well as on the methods of their design.

Lesmo, L., L. Saitta, and P. Torasso, "Fuzzy Production Rules: A Learning Methodology," in P. P. Wang (Ed.), Advances in Fuzzy Sets, Possibility Theory, and Applications, New York: Plenum Press, 1983.

Negoita, C. V., "On Fuzzy Systems," in M. M. Gupta, R. K. Ragade, and R. R. Yager (Eds.), Advances in Fuzzy Set Theory and Applications, Amsterdam, Netherlands: North Holland Publishing Company, 1979.

Negoita, C. V., Expert Systems and Fuzzy Systems, Menlo Park, CA: The Benjamin/Cummings Publishing Company, Inc., 1985.

Nilsson, N. J., "Probabilistic Logic," Artificial Intelligence, Vol. 28, No. 1, 1986.

Prade, H., "Fuzzy Programming: Why and How? - Some Hints and Examples," in P. P. Wang (Ed.), Advances in Fuzzy Sets, Possibility Theory, and Applications, New York: Plenum Press, 1983.

Prade, H., "A Computational Approach to Approximate and Plausible Reasoning With Applications to Expert Systems," IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 7, No. 3, 1985.

The purpose of this paper is twofold: proposing a common basis for the modeling of uncertainty and imprecision, and discussing various kinds of approximate plausible reasoning schemes in this framework. Together with probability, different kinds of uncertainty measures (credibility and plausibility functions in the sense of Shafer, possibility measures in the sense of Zadeh) are introduced in a unified way. The modeling of imprecision in terms of possibility distribution is then presented, and related questions such as the measure of the uncertainty of fuzzy events, the probability and possibility qualification of statements, the concept of a degree of truth, and the truth qualification of propositions, are discussed at length. Deductive inferences from premises weighted by different kinds of measures by uncertainty, or by truth-values in the framework of various multivalued logics, is fully investigated. Then, deductive inferences from imprecise or fuzzy premises are dealt with; patterns of reasoning where both uncertainty and imprecision are present are also addressed. The last section is devoted to the combination of uncertain or imprecise pieces of information given by different sources. On the whole, this paper is a tentative survey of quantitative approaches in the modeling of uncertainty and imprecision, including recent theoretical proposals as well as more empirical techniques such as the ones developed in expert systems as MYCIN or PROSPECTOR, the management of uncertainty and imprecision in reasoning patterns being a key issue in artificial intelligence.

Ralescu, D., "A Survey of the Representation of Fuzzy Concepts and its Applications," in M. M. Gupta, R. K. Ragade, and R. R. Yager (Eds.), Advances in Fuzzy Set Theory and Applications, Amsterdam, Netherlands: North Holland Publishing Company, 1979.

Sanchez, E., "Medical Diagnosis and Composite Fuzzy Relations," in M. M. Gupta, R. K. Ragade, and R. R. Yager (Eds.), Advances in Fuzzy Set Theory and Applications, Amsterdam, Netherlands: North Holland Publishing Company, 1979.

Schefe, P., "On Foundations of Reasoning with Uncertain Facts and Vague Concepts," International Journal of Man-Machine Studies, Vol. 12, 1980.

Shafer, G., "A Theory of Statistical Evidence," in W. L. Harper and C. A. Hooker (Eds.), Foundations and Philosophy of Statistical Theories in Physical Sciences, Hingham, MA: Kluwer Academic, 1975.

Shafer, G., A Mathematical Theory of Evidence, Princeton, NJ: Princeton University Press, 1976.

This book constructs a new theory of epistemic probability. The theory draws on the work of A. P. Dempster but diverges from Dempster's viewpoint by identifying his lower probabilities as epistemic probabilities and taking his rule for combining upper and lower probabilities as fundamental. The book opens with a critique of the well-known Bayesian theory of epistemic probability. It then proceeds to develop an alternative to the additive set functions and the rule of conditioning of the Bayesian theory: set functions that need only be what Choquet called monotone of order $\infty$ , and Dempster's rule for combining such set functions. This rule, together with the idea of weights of evidence, leads to both an extensive new theory and a better understanding of the Bayesian theory. The book concludes with a brief treatment of statistical inference and a discussion of the limitations of epistemic probability. Appendices contain mathematical proofs, which are relatively elementary and seldom depend on mathematics more advanced than the binomial theorem.

Shortliffe, E. H., and Buchanan, B. G., "A Model of Inexact Reasoning in Medicine," Mathematical Biosciences, Vol. 23, 1975.

Medical science often suffers from having so few data and so much imperfect knowledge that a rigorous probabilistic analysis, the ideal standard by which to judge the rationality of a physician's decision, is seldom possible. Physicians nevertheless seem to have developed an ill-defined mechanism for reaching decisions despite a lack of formal knowledge regarding the interrelationships of all the variables that they are considering. This report proposes a quantification scheme which attempts to model the inexact reasoning processes of medical experts. The numerical conventions provide what is essentially an approximation to conditional probability, but offer advantages over Bayesian analysis when they are utilized in a rule-based computer diagnostic system. One such system, a clinical consultation program named MYCIN, is described in the context of the proposed model of inexact reasoning.

Szolovitz, P., and S. G. Pauker, "Categorical and Probabilistic Reasoning in Medical Diagnosis," Artificial Intelligence, Vol. 11, 1978.

Tsotsos, J. K., "On Classifying Time-Varying Events," Proceedings of the IEEE Computer Society Conference on Pattern Recognition and Image Processing, New York: IEEE, 1981.

Whalen, T., and B. Schott, "Decision Support with Fuzzy Production Systems," in P. P. Wang (Ed.), Advances in Fuzzy Sets, Possibility Theory, and Applications, New York: Plenum Press, 1983.

Whalen, T., "Fuzzy Knowledge-Based Systems in Management," Human Systems Management, Vol. 4, No. 4, Autumn 1984.

Knowledge-based systems can be an important complement to management decision making because of their versatility and power in handling partially-structured problems. The chief bottleneck in developing knowledge-based systems is acquiring large amounts of knowledge from human experts. Very closely related is the problem of representing the knowledge in computer-executable form: A good knowledge-representing system can make knowledge acquisition much easier. Using fuzzy production rules with linguistic variables to represent knowledge affords several advantages. Linguistic variables allow human-computer communication in a form resembling a dialogue. Until recently it has seemed to be a gleam in the eyes of academics and science fiction writers. However, recent advances in both hardware and software development means that it is now possible to develop computers which include aspects of intelligence. This article looks at two key areas of advance in terms of what has been achieved and how we can go forward to real applications. The two areas are natural language and intelligent knowledge-based systems.

Yager, R. R., "A New Methodology for Ordinal Multiobjective Decisions Based on Fuzzy Sets," Decision Sciences, Vol. 12, 1981.

Using the structural forms supplied by fuzzy set theory and approximate reasoning, a new method is presented for solving multiple-objective decision problems for which the decision maker can supply only ordinal information on his preferences and the importance of the individual objectives.

Zadeh, L. A., "Fuzzy Sets," Information and Control, Vol. 8, 1965.

A fuzzy set is a class of objects with a continuum of grades of membership. Such a set is characterized by a membership (characteristic) function which assigns to each object a grade of membership ranging between zero and one. The notions of inclusion, union, intersection, complement, relation, convexity, etc., are extended to such sets, and various properties of these notions in the context of fuzzy sets are established. In particular, a separation theorem for convex fuzzy sets is proved without requiring that the fuzzy sets be disjoint.

Zadeh, L. A., "Fuzzy Logic and Approximate Thinking," Synthese, Vol. 30, 1975.

Zadeh, L. A., "Possibility Theory as a Basis for Information Processing and Knowledge Representation," Proceedings of the COMPSAC 80, Oct. 1980.

To enhance the effectiveness of computers as information processing machines it is necessary to develop a better understanding of ways of dealing with information which is imprecise, incomplete or not totally reliable. The recently developed theory of possibility provides a non-statistical basis for the representation and manipulation of uncertain information. The central concept in this theory is that of a possibility distribution, by which is meant the fuzzy set of all possible values of a variable. With this concept as the point of departure, one can develop a set of translation rules for finding the possibility distribution which is induced by a given proposition. Such rules form a part of the representation language, PRUF, which in combination with fuzzy logic, provides a formal system for the representation and inference from imprecise knowledge.

Zadeh, L. A., "Possibility Theory and Soft Data Analysis," Proceedings of the AAAS Selected Symposium 54 on Matehmatical Frontiers of the Social and Policy Sciences, Washington D.C.: American Association for the Advancement of Science, 1981.

A thesis advanced in this paper is that much of the uncertainty which is associated with soft data is nonstatistical in nature. Based on this premise, an approach to the representation and manipulation of soft data -- in which the recently developed theory of possibility plays a central role -- is described and illustrated with examples.

Zadeh, L. A., "Fuzzy Probabilities and Their Role in Decision Analysis," Proceedings of the Fourth MIT/ONR Workshop on Distributed Information and Decision Systems, 1982.

The conventional approaches to decision analysis are based on the assumption that the probabilities which enter into the assessment of the consequences of a decision are known numbers. In most realistic settings, this assumption is of questionable validity since the data from which the probabilities must be estimated are usually incomplete, imprecise or not totally reliable. In the approach outlined in this paper, the probabilities are assumed to be fuzzy rather than real numbers. It is shown how such probabilities may be estimated from fuzzy data and a basic relation between joint, conditional and marginal fuzzy probabilities is established. Manipulation of fuzzy probabilities requires, in general, the use of fuzzy arithmetic, and many of the properties of fuzzy probabilities are simple generalizations of the corresponding properties of real-valued probabilities.

Zadeh, L. A., "The Role of Fuzzy Logic in the Management of Uncertainty in Expert Systems," Fuzzy Sets and Systems, Vol. 11, 1983a.

Management of uncertainty is an intrinsically important issue in the design of expert systems because much of the information in the knowledge base of a typical expert system is imprecise, incomplete, or not totally reliable. In the existing expert systems, uncertainty is dealt with through a combination of predicate logic and probability-based methods. A serious shortcoming of these methods is that they are not capable of coming to grips with the pervasive fuzziness of information in the knowledge base, and, as a result, are mostly ad hoc in nature. An alternative approach to the management of uncertainty, which is suggested in this paper, is based on the use of fuzzy logic, the logic underlying approximate or, equivalently, fuzzy reasoning. A feature of fuzzy logic, which is of particular importance to the management of uncertainty in expert systems, is that it provides a systematic framework for dealing with fuzzy quantifiers (e.g., most, many, few, not very many, almost all, infrequently, about 0.8, etc.). In this way, fuzzy logic subsumes both predicate logic and probability theory, and makes it possible to deal with different types of uncertainty within a single conceptual framework. In fuzzy logic, the deduction of a conclusion from a set of premises is reduced, in general, to the solution of a nonlinear program through the application of projection and extension principles. This approach to deduction leads to various basic syllogisms which may be used as rules of combination of evidence in expert systems. Among syllogisms of this type, which are discussed in this paper, are the intersection/product syllogism, the generalized modus ponens, the consequent conjunction syllogism, and the major-premise reversibility rule.

Zadeh, L. A., "A Theory of Commonsense Knowledge," Memorandum No. UCB/ERL M83/26, Electronics Research Laboratory, College of Engineering, University of California, Berkeley, CA, 1983b.

The theory outlined in this paper is based on the idea that what is commonly called commonsense knowledge may be viewed as a collection of dispositions, that is, propositions with implied fuzzy quantifiers. Typical examples of dispositions are: Icy roads are slippery, Tall men are not very agile, Overeating causes obesity, Bob loves women, What is rare is expensive, etc. It is understood that, upon restoration of fuzzy quantifiers, a disposition is converted into a proposition with explicit fuzzy quantifiers (e.g., tall men are not very agile; most tall men are not very agile). Since traditional logical systems provide no methods for representing the meaning of propositions containing fuzzy quantifiers, such systems are unsuitable for dealing with commonsense knowledge. It is suggested in this paper that an appropriate computational framework for dealing with commonsense knowledge is provided by fuzzy logic, which, as its name implies, is the logic underlying fuzzy (or approximate) reasoning. Such a framework, with an emphasis on the representation of dispositions, is outlined and illustrated with examples.

Zadeh, L. A., "Linguistic Variables, Approximate Reasoning and Dispositions," Medical Information, Vol. 8, No. 3, 1983c.

Zadeh, L. A., A Computational Approach to Fuzzy Quantifiers in Natural Languages," Computer and Mathematics with Applications, Vol. 9, No. 1, 1983d.

The generic term, fuzzy quantifier, is employed in this paper to denote the collection of quantifiers in natural languages whose representative elements are: several, most, much, not many, very many, not very many, few, quite a few, large number, small number, close to five, approximately ten, frequently, etc. In our approach, such quantifiers are treated as fuzzy numbers which may be manipulated through the use of fuzzy arithmetic and, more generally, fuzzy logic. A concept which plays an essential role in the treatment of fuzzy quantifiers is that of the cardinality of a fuzzy set. Through the use of this concept, the meaning of a proposition containing one or more fuzzy quantifiers may be represented as a system of elastic constraints whose domain is a collection of fuzzy relations in a relational database. This representation, then, provides a basis for inference from premises which contain fuzzy quantifiers. For example, from the propositions "Most U's are A's" and "Most A's are B's", it follows that "Most U's are B's," where most is the fuzzy product of the fuzzy proposition most with itself. The computational approach to fuzzy quantifiers which is described in this paper may be viewed as a derivative of fuzzy logic and test-score semantics. In this semantics, the meaning of a semantic entity is represented as a procedure which tests, scores and aggregates the elastic constraints which are induced by the entity in question.

Zadeh, L. A., "Fuzzy Sets and Commonsense Reasoning" Berkeley Cognitive Science Report No. 21, Cognitive Science Program, Institute of Human Learning, University of California, Berkeley, CA, 1984a.

Despite the extensive literature and numerous conferences, there are many misconceptions about the concept of a fuzzy set which find their expression in questions such as: What is a fuzzy set? Is there a difference between fuzziness and randomness? Are there significant applications in which fuzzy set techniques have a distinct advantage over conventional probabilistic methods? What is the role that the theory of fuzzy sets can play in artificial intelligence, expert systems, operations research, process control, medical diagnosis and related fields? These are some of the questions that the author attempts to answer -- directly or indirectly -- in this article.

Zadeh, L. A., Decision Analysis and Fuzzy Mathematics, Berkeley, CA: Computer Science Division, University of California, 1984b.

In sum, the concept of a linguistic variable, in combination with the other basic concepts of fuzzy mathematics which derive from the notion of a fuzzy set, make it possible to formulate and solve a number of problems in decision analysis that are too complex or too ill-defined to be susceptible to solution by conventional techniques. In this way, the theory can be brought closer to reality. There will always be many problems, however, for which no theory, whether based on fuzzy mathematics or not, could provide a rational basis for solution. This is particularly true of problems relating to conflict resolution, group decision analysis, and, more generally, to large-scale systems in which the goals and horizons are hard to agree upon and not merely hard to quantify. Thus, in the final analysis, the ultimate limitations of any theory of decision processes derives not from its mathematical foundations but from our inability to understand and elucidate the meaning of the rationality of human behavior. This will always be a controversial issue.

Zadeh, L. A.,"Management of Uncertainty in Expert Systems," Proceedings of the First Symposium on the Theory and Application of Expert Systems in Emergency Management Operations, Special Publication 717, Gaithersburg, MD: National Bureau of Standards, 1986.

Zimmerman, H. J., L. A. Zadeh, and B. R. Gaines (Eds.), "Fuzzy Sets and Decision Analysis," TIMS Studies in the Management Sciences, Vol. 20, Amsterdam, Netherlands: North-Holland Publishing Company, 1984.

# 8. APPLICATIONS OF EXPERT SYSTEMS

## 8.1 TYPES OF PROBLEMS AMENABLE TO SOLUTIONS BY EXPERT SYSTEMS

Recalling the definition of expert systems - a computer program capable of solving those decision problems that normally require human expertise - we see that the range of such problems can be rather extensive. A distinction can be made, however, between those decision problems that have well-documented procedures or algorithms for solving them and those that do not. The latter set of problems is the domain of expert systems.

The problems in this domain are further characterized by incomplete data and uncertainty regarding the information available to solve them. Experts, using their experience, represented by rules-of-thumb (heuristics) and subjective judgments, are able to fill these gaps and, thus, provide solutions. If expert systems are to solve these decision problems, then they must mimic the human expert by using the same techniques as the human expert and/or apply other decision procedures that prove to produce comparable results. Thus, we would expect expert systems to be able to represent and manipulate data, deal with incomplete and uncertain data, and use heuristic rules and make subjective judgments. To form a practical expert system, it is necessary to extract the knowledge and methodologies of the human expert; this may not be possible for most problems. As the art of building expert systems stands today, it is only possible to resolve those problems that have a narrow domain of expertise.

There are many criteria that may be used to decide whether or not to build an expert system and for selecting suitable problems. First, if an expert system is to be worth the effort, its performance must be at least comparable to that of a human expert in that field. Second, the speed of response may be important. In such instances, one should be able to take advantage of an expert system's unique capacity to combine its human expert-like capabilities with its already existing computing speed. Third, there are many complex and urgent problems requiring responses when experts are not at hand. Here, expert systems can be used in aiding the non-experts to make the necessary responses. Fourth, there are certain problem areas where even the experts wish they had an intelligent assistant who can be a source of information and ideas, and with whom they could interact in making decisions. An expert system can be used for these purposes. Finally, for many problems, there is an acute shortage of experts, and new ones cannot be trained very quickly, see Basden [1983] for further discussion.

## 8.2 DOMAINS IN WHICH EXPERT SYSTEMS ARE BEING CURRENTLY APPLIED

Over the past decade or so, expert systems have been built for many problem domains, the majority of them for problems requiring diagnostic capabilities. They differ in the type of global data base used, the makeup of their inference engine, as well as in the type of knowledge representations required to make their decisions (see Chapter 3 for an explanation of these concepts). Tables 8.1 through 8.7 provide a breakdown of areas into specific problem domains for which expert systems have been built. Readers interested in the details about a particular system are directed to the cited references in Tables 8.1 through 8.7.

## Applications in Medicine

Expert systems in the medical field were some of the earliest (and the most successful) ones to be built. The majority of these systems aid in the diagnosis of diseases in a narrowly defined domain of medicine. Many are also capable of recommending treatments. For example, CASNET aids in the diagnosis, as well as in the treatment of glaucoma. Table 8.1 lists a number of medical expert systems, along with their areas of expertise.

TABLE 8.1 Expert Systems in Medicine

| Expert Systems | Problem Domain/Source Article |
|---|---|
| CASNET | Diagnosis and therapy recommendations for glaucoma<br>Kulikowski and Weiss [1982] |
| CENTAUR | Interpretation of pulmonary function test measurements from patients with lung disorders<br>Aikens [1983] |
| DIGITALIS ADVISOR | Treatment recommendation for using digitalis<br>Swartout [1977] and Gorry, Silverman and Pauker [1978] |
| EMERGE | Analysis of chest pain in the emergency room<br>Hudson and Estrin [1984] |
| EMYCIN | A domain-independent version of MYCIN, usable for developing rule-based consultation programs for many fields<br>van Melle [1979] |
| HEADMED | Diagnostic advice in psychopharmacology<br>Heiser, Brooks and Ballard [1978] |
| INTERNIST/CADUCEUS | Diagnosis in internal medicine<br>Pople [1973] and Pople [1982] |
| KARDIO-E | Diagnosis of cardiac arrhythmias using electrocardiographs<br>Lavroc [1985] |
| KMS | Diagnosis of medical problems<br>Reggia et al. [1980] |

TABLE 8.1 Expert Systems in Medicine (Continued)

| Expert System | Problem Domain/Source Article |
|---|---|
| MDX | Diagnosis of causes for cholestasis<br>Chandrasekaran and Mittal [1983] |
| MYCIN | Diagnosis and treatment recommendations for certain bacterial infections<br>Shortliffe [1976] |
| ONCOCIN | Assistance in the management of cancer patients on chemotherapy protocols for forms of lymphoma<br>Shortliffe et al. [1981] |
| PATREC | Intelligent medical data base needed for MDX<br>Mittal, Chandrasekaran and Stickler [1984] |
| PIP | Clinical diagnosis<br>Parker et al. [1976] |
| PUFF | Analysis of data from pulmonary function tests<br>Atkins et al. [1982] |
| RADEX | Diagnostic advice in radiology through interpretation of various kinds of imaging data<br>Chandrasekaran, Mittal, and Smith [1983] |
| RX | Determination of causal relationships in medicine<br>Blum [1982] |
| TIERESIAS | Generation of explanations of MYCIN's diagnoses and treatment recommendations<br>Davis [1976, 1982] |
| VM | Monitoring of patient's respiration during intensive care<br>Fagan [1980] |

## Applications in Chemistry

Many of the problem domains in chemistry also require diagnostic skills. However, unlike medical diagnosis, where medical data and evidence are generally used to draw conclusions that already exist in the knowledge base, many chemistry expert systems generate new possibilities for the user to examine. This capability is extremely useful to the chemist trying to identify the molecular structure of a chemical compound. Table 8.2 gives a representation of the range of problems in the area of chemistry for which expert systems have been built.

TABLE 8.2 Expert Systems in Chemistry

| Expert System | Problem Domain/Source Article |
|---|---|
| CONGEN | Construction of chemical structures<br>Cahart [1979] |
| CRYSALIS | Identification of structure of a protein from electron-density maps<br>Engelmore and Nii [1977] |
| DENDRAL/METADENDRAL | Generation of structural representations of organic molecules using mass spectroscopic data for their identification<br>Buchanan and Feigenbaum [1978] |
| GA1 | Inference of a complete molecular structure by generating molecular pieces<br>Stefik [1978] |
| GAMMA | Interpretation of gamma ray activation spectra<br>Barstow [1979] |
| MOLGEN | Designing of molecular genetic experiments<br>Friedland [1980] and Martin et al. [1977] |
| SECS | Generation of schemes for synthesis of organic compounds<br>Wipke et al. [1977] |
| SYNCHEM | Generation of schemes for synthesis of organic compounds<br>Gelernter et al. [1977] |

## Applications in Engineering, Mathematics, and other Sciences

The expert systems in engineering, mathematics and other sciences have been built to solve difficult problems requiring complex representation and manipulation. Many of the applications attempted in the engineering domain are similar to those in chemistry -- structural analysis, circuit analysis and design, etc. Some of the better-known expert systems in these areas are given in Table 8.3.

TABLE 8.3 Expert Systems in Engineering, Mathematics, and Other Sciences

| Expert System | Problem Domain/Source Article |
|---|---|
| AM | Discovery of mathematical concepts<br>Lenat [1982] |
| ASA | Analysis of experiments using non-parametric statistical tests<br>O'Keefe [1982] |
| EL | Analysis of electrical circuits<br>Sussman and Stallman [1975] |
| EURISKO | Generation of solutions through the use of heuristics<br>Lenat [1983] |
| MACSYMA | Symbolic mathematics<br>Moses [1971] |
| MECHO | Problem-solving in mechanics<br>Bundy et al. [1979] |
| REX | Performance of regression analysis, interpretation of results, and tutoring<br>Gale and Pregibon [1982] |
| SACON | Advice in the use of large, general-purpose structural analysis programs to engineers<br>Bennett and Engelmore [1979] |
| STATPATH | Identification of a statistical analysis technique appropriate for the data in hand<br>Portier and Lai [1983] |
| SU/X | Understanding signals (represented as descriptions of spectral lines) for identification, location, and determination of velocity of objects in a physical space<br>Nii and Feigenbaum [1978] |

## Applications in Oil/Mineral Exploration

This is another area in which expert system methodology has been successfully applied. Here, the expert system serves the same purpose as an expert by using the same data and reasoning to analyze and predict the chances of successful oil strikes or mineral finds. Some of the better known expert systems in this area are given in Table 8.4.

TABLE 8.4 Expert Systems in Oil/Mineral Exploration

| Expert System | Problem Domain/Source Article |
|---|---|
| DIPMETER ADVISOR | Analysis of dipmeter logs to determine the subsurface tilt in oil drilling operations<br>Davis et al. [1981] |
| DRILLING ADVISOR | Determination of cause and possible resolution of problems encountered in the course of drilling production oil wells<br>Hollander et al. [1983] |
| ELAS | Well-log (various electromagnetic, sonic, and nuclear signals) analysis through multilevel reasoning<br>Weiss et al. [1982] |
| KAS | Acquisition of knowledge needed for PROSPECTOR<br>Reboh [1982] |
| PROSPECTOR | Aid in mineral exploration<br>Gaschnig [1982] and Duda, Gaschnig and Hart [1977] |

## Applications in Military and Government

Applications in these areas represent some of the more practical uses for which expert systems have been applied. In many of these systems, the expert system technology is often combined with other tools to process the necessary data. Efficiency and speed have been the most important gains in using expert systems in these problem domains. Many of the systems listed in Table 8.5 are still in the testing stages.

TABLE 8.5 Expert Sytems in Military and Government

| Expert System | Problem Domain/Source Article |
|---|---|
| AIRID | Aircraft identification<br>Aldridge [1984] |
| AIRPLAN | Assistance in launch and recovery of aircraft on aircraft carriers<br>Masui, McDermott and Sobel [1983] |
| BATTLE | Battlefield weapons assignment<br>Slagle et al. [1984] |
| HASP/SIAP | Ocean surveillance through signal-processing<br>Nii et al. [1982] |
| KNOBS | Tactical mission planning consultant<br>Engelman, Berg and Bischoff [1979] |
| SPAM | Interpretation of high resolution aerial photographs of commercial airports<br>McKeown and McDermott [1983] |
| SPERIL | Earthquake damage assessment for structures<br>Ishiguka, Fu and Yao [1981] |
| STAMMER2 | Tactical situation assessment<br>McColl [1979] |
| TAC II | Tactical Decision Making<br>Geschke, Bullock and Widmaier [1983] |
| TATR | Tactical targeting<br>Callero, Jamison and Waterman [1982] |

## Applications in Industry and Business

Industry and business seems to have been the main beneficiary of expert systems, if the number of application areas in which they have used this technology is any indication. One of the most frequently cited success stories is that of R1 (also called XCON) used by the Digital Electronics Corporation to configure its computers. Table 8.6 lists problem domains for which expert systems have been constructed.

TABLE 8.6 Expert Systems in Industry and Business

| Expert System | Problem Domain/Source Article |
|---|---|
| ACE | Generation of trouble-shooting reports and management analyses for telephone cable maintenance<br>Vesonder, Salvatore and Zielinski [1983] |
| ACRONYM | Visual perception through analysis of photographs<br>Binford [1981] |
| APRIKS | Diagnosis and treatment of plant diseases, and crop pest control<br>Tou and Cheng [1983] |
| AQ11 | Diagnosis of plant diseases<br>Michalski and Chilansky [1980] |
| CATS-1 | Diagnosis of diesel locomotive malfunctions and demonstration of repair procedures<br>Kinnucan [1984] |
| CRIB | Computer fault-finding through knowledge engineering<br>Hartley [1984] |
| DART/DASD | Diagnosis of computer system failures<br>Bennett and Hollander [1981] |
| GETREE | Knowledge management through interactive graphics in building expert systems<br>Lewis and Lynch [1983] |
| HARPY | Speech understanding<br>Newell [1978] |
| HEARSAY | Speech understanding<br>Hayes-Roth and Lesser [1977] |

TABLE 8.6 Expert Systems in Industry and Business (Continued)

| Expert System | Problem Domain/Source Article |
|---|---|
| NDS | Isolation and repair of multiple faults in electronic systems<br>Williams, Orgren and Smith [1983] |
| NOAH | Use of robot in problem solving<br>Sacerdoti [1974] |
| PTRANS | Assistance in processing customer orders for computer systems<br>McDermott [1983] |
| R1 (XCON) | Configuration of computer systems based on customer orders for components<br>Shannon [1984] |
| REACTOR | Diagnosis and treatment of nuclear reactor accidents<br>Nelson [1982] |
| STEAMER | Determination of causes and remedies for mechanical problems<br>Forbus [1981] and Stevens et al. [1981] |
| STRIPS | Use of robot in problem solving<br>Nilsson [1980] |
| VENTILATION EXPERT | Management of the operation of coal mine ventilation systems<br>Kohler [1986] |
| XSEL | Assistance to salespeople in tailoring computer systems to fit the needs of the customer<br>McDermott [1982] |

## Applications in Education and Training

The earliest systems in this domain were simple programs with limited tutoring capabilities. Today's expert systems not only tutor the students on how to solve complex problems, but also assess their understanding and adjust the tutoring approach based on the student's sophistication. Many successful expert systems in this area are really descendents of expert systems originally built to perform certain specific tasks. Table 8.7 gives the names and domains of some expert systems in this area.

TABLE 8.7 Expert Systems in Education and Training

| Expert System | Problem Domain/Source Article |
|---|---|
| ACE | Environment allowing students to explore algorithms<br>Miller et al. [1983] |
| BUGGY | Determining a student's arithmetic misconceptions<br>Brown and Burton [1982] |
| EXAMINER | Tutoring in the use of contextual knowledge in the analysis of diagnostic behavior<br>Olesson [1977] |
| GUIDON | Tutoring through case-method for improving a student's ability to diagnose complex problems in medicine and science<br>Clancey [1979] |
| LOGO | Learning environment for student initiated learning of a broad range of subjects<br>Howe and O'Shea [1978] |
| NEOMYCIN | Interpretation of student behavior and teaching diagnostic strategy in medicine<br>Clancey and Letsinger [1981] |
| RESEDA | Biographical data management<br>Zarri [1984] |
| SCHOLAR | Tutoring in geography of South America<br>Carbonell[1970] |
| SOPHIE | Tutoring in electronic trouble shooting<br>Brown, Burton and de Kleer [1982] |
| WUMPUS | Tutoring in logical and reasoning skills<br>Goldstein [1982] |

## REFERENCES


Addis, T. R., "Towards an 'Expert' Diagnostic System," ICL Technical Journal, 1980.

Aikins, J. S., "Prototypical Knowledge for Expert Systems," Artificial Intelligence, Volume 20, No. 2, February 1983.

Aikins, J. S., et al., "PUFF: An Expert System for Interpretation of Pulmonary Function Data," Report No.: STAN-HPP-82-13, Stanford University, Computer Science Department, Stanford, CA, 1982.

Aldridge, J. P., "AIRID: An Application of the KAS/Prospector Expert System Builder to Airplane Identification," Report No.: LA-UR-84-988; CONF-8404128-1, Los Alamos National Lab., NM, 1984.

The Knowledge Acquisition System/Prospector expert system building tool developed by SRI, International, has been used to construct an expert system to identify aircraft on the basis of observables such as wing shape, engine number/location, fuselage shape, and tail assembly shape. Additional detailed features are allowed to influence the identification as other favorable features. Constraints on the observations imposed by bad weather and distant observations have been included as contexts to the models. Models for Soviet and US fighter aircraft have been included. Inclusion of other types of aircraft such as bombers, transports, and reconnaissance craft is straightforward. Two models permit exploration of the interaction of semantic and taxonomic networks with the models. A full set of text data for communication with the user has been included. This paper presents discussion of the ease of building the expert system using this powerful tool and problems encountered in the construction process.

Barstow, D., "An Experiment in Knowledge-Based Automatic Programming," Artificial Intelligence, Vol. 12, 1979.

Barstow, D., "Knowledge engineering in Nuclear Physics," Proceedings of the Sixth IJCAI, American Association for Artificial Intelligence, 1979.

Basden, A., "On the Application of Expert Systems," International Journal of Man-Machine Studies, Vol. 19, 1983.

This article seeks to bring together a number of issues relevant to the application of expert systems by discussing their advantages and limitations, their roles and benefits, and the influence that real-life applications might have on the design of expert systems software. Part of the expert systems strategy of one major chemical company is outlined. This system is described briefly and used to illustrate much of the later discussion. It is of the plausible-inference type and has application in the field of materials engineering. The article is aimed both at the interested end-user and those working in the field of expert systems.

Basden, A., and B. A. Kelly, "An Application of Expert Systems Techniques in Materials Engineering," Proceedings of the Colloquium on Application of Knowledge-Based (or Expert) Systems, London, England, 1982.

Bennett, J. and R. Engelmore, "SACON: A Knowledge-Based Consultant for Structural Analysis," Proceedings of the Sixth IJCAI, American Association for Artificial Intelligence, 1979.

This paper presents an application of Artificial Intelligence methods to the engineering domain of structural analysis. The authors have developed and partially implemented an "automated consultant" called SACON (Structural Analysis CONsultant), using the EMYCIN system as its frame work. SACON advises engineers in the use of a large, general-purpose structural analysis program. The structure of the knowledge base, including the major concepts used and inferences drawn by consultants, is presented. Some observations in the light of this application about the EMYCIN system as a representational vehicle are given.

Bennett, J. S., and C. R. Hollander, "DART: An Expert System for Computer Fault Diagnosis," Proceedings of the Seventh IJCAI, American Association for Artificial Intelligence, 1981.

This paper describes an application of artificial intelligence techniques to computer system fault diagnosis. In particular, the authors have implemented an automated consultant that advises IBM field service personnel on the diagnosis of faults occurring in computer installations. The consultant identifies specific system components (both hardware and software) likely to be responsible for an observed fault and offers a brief explanation of the major factors and evidence supporting these indictments. The consultant, called DART, was constructed using EMYCIN, and is part of a large research effort investigating automated diagnosis of machine faults.

Bennett, S. W. and A. C. Scott, "Computer-Assisted Customized Antimicrobial Dosages," American Journal of Hospital Pharmacy, Vol. 37, 1980.

Blidberg, D. R., A. S. Westneat, and R. W. Corell, "Expert Systems, a Tool for Autonomous Underwater Vehicles," Proceedings of Trends and Applications 1983 on Automating Intelligent Behavior: Applications and Frontiers, Silver Spring, MD: IEEE Computer Society Press, May 1983.

The need for guidance and control of an unmanned untethered underwater automaton calls for a knowledge-based expert as pilot. This paper considers the expert system in a time-constrained environment. It concludes that substantial improvements in reliability, in performance of the mission, in coping with unexpected events should be achievable with on-board microcomputer capabilities.

Blum, R. L., "Discovery, Confirmation, and Incorporation of Causal Relationships From a Large Time-Oriented Clinical Data Base: The RX Project," Computers and Biomedical Research, Volume 15, No. 2, 1982.

Brodsky, S., and N. Tyle, "Knowledge-Based Expert Systems for Power Engineering," Proceedings of the Fifteenth Pittsburgh Modeling and Simulation Conference, Pittsburgh, PA, Aug. 1984.

Brooke, J. B., A. L. Rector, and M. G. Sheldon, "A Review of Studies of Decision-Making in General Practice," Medical Information, Vol. 9, No. 1, 1984.

Recent advances in expert systems and statistics make the development of decision-support systems for clinical use a real possibility. One natural location for such systems is with the general practitioner, as the general practitioner has to manage a wide range of problems covering all specialties, and often has to make decisions without the necessary knowledge at hand. However, if decision-support systems are to gain wide acceptance, a number of criteria must be met. They must be easy to use and not interfere with the consultation. More important, they must provide information about important problems which clinicians actually face in their consulting rooms, and they must provide help in ways which are compatible with the way doctors think and work. The investigation of decision-making in general practice is therefore one of the necessary prerequisites of any successful developments of the next generation of computer systems for clinical use.

Brown, J. S., and R. R. Burton, "Diagnosing Bugs in a Simple Procedural Skill," in D. Sleeman and J. S. Brown (Eds.), Intelligent Tutoring Systems, New York: Academic Press, 1982.

Brown, J. S., R. R. Burton, and J. de Kleer, "Pedagogical, Natural Language and Knowledge Engineering Techniques in SOPHIE I, II and III," in D. Sleeman and J. S. Brown (Eds.), Intelligent Tutoring Systems, New York: Academic Press, 1982.

Buchanan, B. G., and E. A. Fiegenbaum, "DENDRAL and Meta-DENDRAL: Their Applications Dimension," Artificial Intelligence, Vol. 11, 1978.

The DENDRAL and Meta-DENDRAL programs are products of a large interdisciplinary group of Stanford University scientists concerned with many and highly varied aspects of the mechanization of scientific reasoning and the formalization of scientific knowledge for this purpose. An early motivation for our work was to explore the power of existing AI methods, such as heuristic search, for reasoning in difficult scientific problems. Another concern has been to exploit the AI methodology to understand better some fundamental questions in the philosophy of science, for example, the process by which explanatory hypotheses are discovered or judged adequate. From the start, the project has had an application dimension. It has sought to develop "expert level" agents to assist in the solution of problems in their discipline that require complex symbolic reasoning. The applications dimension is the focus of this paper. In order to achieve high performance, the DENDRAL programs incorporate large amounts of knowledge about the area of science to which they are applied, structure elucidation in organic chemistry. A "smart assistant" for a chemist needs to be able to perform many tasks, as does an expert, but need not necessarily understand the domain at the same theoretical level as the expert. The over-all structure elucidation task is described followed by a description of the role of the DENDRAL programs within that framework. The Meta-DENDRAL programs use a weaker body of knowledge about the domain of mass spectrometry because their task is to formulate rules of mass spectrometry by induction from empirical data.

Cahart, R. E., "CONGEN: An Expert System Aiding the Structural Chemist," in D. Michie (Ed.), Expert Systems in the Micro-Electronic Age, Edinburgh, Scotland: Edinburgh University Press, 1979.

CONGEN (for CONstrained GENerator) is a complex symbol-manipulation program for the construction of chemical structures. Its input consists of a molecular formula (i.e., a list of atoms to be used as building blocks) together with restrictions (constraints) on the possible interconnections among the atoms. Using a variety of sophisticated graph-theoretic algorithms, CONGEN finds all possible ways of assembling the atoms into molecular structures which satisfy the specified constraints. Aside from a description of CONGEN and some examples of its contributions to structure analysis, the discussion focuses upon two issues which are important in any application of expert systems. One relates to the amount of human engineering needed to overcome the psychological barriers to the use of complex computer programs. This issue not only encompasses the design of a smooth and helpful user-interface, but also of basic algorithms to reformulate the user's input. There are often several alternative ways of expressing a problem to a complex expert system, and what seems to the user to be a perfectly logical statement of the problem may be far from optimum from the program's standpoint. The other issue involves the choice between different representations of the expert knowledge contained in the sytems: What parts of the knowledge should be hand coded into the system as highly efficient, special-purpose algorithms, and what parts should be left in a more flexible, external form?

Callero, M. L. Jamison, and D. A. Waterman, "TATR: An Expert Aid for Tactical Air Targeting (Interim report)," Report No.: RAND/N-1796-ARPA, RAND Corp., Santa Monica, CA, Jan. 82.

This report describes an initial version of TATR (Tactical Air Target Recommender), a prototype expert system being developed at RAND to assist tactical air targeters in selecting and prioritizing targets. TATR applies a knowledge engineering problem-solving approach in which human domain knowledge is essential, and judgment plays a larger role than mathematical algorithms and stochastic formalisms. Under interactive user direction, TATR preferentially orders enemy airfields, determines targets on those airfields to attack, and identifies the most effective weapons systems against those targets. TATR is programmed in the ROSIE language, that has an English-like syntax that facilitates nonprogrammer comprehension and program verification. The program applies predetermined planning heuristics to generate an airfield attack plan. It then replans, incorporating user modifications, and projects a series of plans over several days. TATR interactively maintains databases by requesting and processing updates from the user and provides detailed information about plans, friendly force capability, and enemy force posture and status.

Campbell, A. N., et al., "Recognition of a Hidden Mineral Deposit by an Artificial Intelligence Program," Science, Vol. 217, No. 3, 1982.

Carbonell, J., "AI and CAI: An Artificial-Intelligence Approach to Computer-Assisted Instruction," IEEE Transactions on Man-Machine Systems, MMS-11, 1970.

Chandrasekaran, B., and S. Mittal, "Conceptual Representation of Medical Knowledge for Diagnosis by Computer: MDX and Related Systems," in Marshall C. Yovits (Ed.), Advances in Computers, Vol. 22, New York: Academic Press, 1983.

The purpose of this article is to describe an approach to the design of medical decision-making systems based on the notion of conceptual structures for knowledge representation. A collection of related systems that have been under development in the authors' laboratory exemplifies this approach, but the areas are more general than the particular systems to be described. Section 2 provides an overview, from a theoretical viewpoint, of the conceptual structure methodology. Later sections describe the functioning of the systems being developed to give concreteness to the theoretical ideas. The central system in this group of systems is called MDX, which is a diagnostic system (i.e., it attempts to classify a given case as an element of a disease taxonomy). This system interacts with two other systems during its problem solving, PATREC and RADEX. The former is a knowledge-based patient data base system that answers MDX's queries about patient data, and the latter is a radiological consultant which helps MDX in the interpretation of various kinds of imaging data. Both PATREC and RADEX are invoked by MDX as needed, but MDX is in control of the overall diagnostic process.

Chandrasekaran, B., S. Mittal, and J. W. Smith, "RADEX: Towards a Computer-Based Radiology Consultant," in E. S. Gelesma and L. N. Kanal (Eds.), Pattern Recognition in Practice, Amsterdam, Netherlands: North-Holland, 1980.

Chang, C. L., "Decision Support in an Imperfect World," Proceedings of Trends and Applications 1983 on Automating Intelligent Behavior: Applications and Frontiers, Silver Spring, MD: IEEE Computer Society Press, May 1983.

Clancey, W. J., "Tutoring Rules for Guiding a Case Method Dialogue," International Journal of Man-Machine Studies, Volume 11, 1979.

Clancey, W. J., and R. Letsinger, "NEOMYCIN: Reconfiguring a Rule-Based Expert System for Application to Teaching," Proceedings of the Seventh IJCAI, American Association for Artificial Intelligence, 1981.

NEOMYCIN is a medical consultant system in which MYCIN's knowledge base is reorganized and extended for use in GUIDON, a teaching program. The new system constitutes a psychological model for doing diagnosis, designed to provide a basis for interpreting student behavior and teaching diagnostic strategy. The model separates out kinds of knowledge that are procedurally embedded in MYCIN's rules and are inaccessible to the teaching program. The key idea is to represent explicitly and separately: a domain-independent diagnostic strategy in the form of meta-rules, knowledge about the structure of the problem space, causal and data/hypothesis rules, and world facts. As a psychological model, NEOMYCIN captures the forward-directed, compiled association mode of reasoning that characterizes expert behavior. Collection and interpretation of data are focused by the differential or working memory hypotheses. Moreover, the knowledge base is broadened so that GUIDON can teach a student when to consider a specific infectious disease and what competing hypotheses to consider (i.e., essentially the knowledge a human would need in order to use the MYCIN consultation system properly).

Clancey, W. J., and E. H. Shortliffe (Eds.), Readings in Medical Artificial Intelligence: The First Decade, Reading, MA: Addison-Wesley Publishing Company, 1984.

Davis, R., "TEIRESIAS: Applications of Meta-Level Knowledge," in R. Davis and D. B. Lenat (Eds.), Knowledge-Based Systems in Artificial Intelligence, New York: McGraw-Hill Book Company, 1982.

Davis, R., et al., "The DIPMETER ADVISOR: Interpretation of Geologic Signals," Proceedings of the Seventh IJCAI, American Association for Artificial Intelligence, 1981.

de Mori, R., et al., "An Expert System for Speech Decoding," Proceedings of the ECAI-82, 1982.

Dillard, J. F., and K. Ramakrishna, "Knowledge-Based Expert Computer Systems for Accounting Related Tasks: The Case of Price Analysis," Proceedings of the Sixteenth Annual Meeting of the American Institute for Decision Sciences, Atlanta, GA: American Institute for Decision Sciences, Nov. 1984.

Artificial intelligence (AI) technology is applied in constructing a prototype system to assist in price analysis decisions. This is accomplished by using indepth interviews, verbal protocols and available authoritative documentation. The resulting knowledge base is used to design the architecture of a prototype, which in turn provides the basic structure for building more intelligent expert computer systems in this and other related decisions situations.

Duda, R. O., and T. D. Garvey, "A Study of Knowledge-Based Systems for Photo Interpretation," Artificial Intelligence Center, SRI International, Menlo Park, CA, 1980.

Dutta, A., and A. Basu, "An Artificial Intelligence Approach to Model Management in Decision Support Systems," Computer, Vol. 17, No. 9, Sept. 1984.

Decision support systems, or DSS's, are computer-based systems that can be used directly by decision makers who are not sophisticated programmers to solve semistructured or unstructured problems. Such problems are characterized by the lack of quantitative descriptions, well-defined goals, or prescribed algorithms for their solution. As a result, the distribution of effort in solving them differs significantly from that required for structured problems. Since many important management problems, especially strategic problems, tend to be unstructured, or semistructured at best, the development of effective DSS's is an active field of research. Generically, a DSS consists of three major subsystems; the dialog, data, and the models. Due to the characteristics of DSS applications, the solution process usually involves appreciable trial-and-error, and data is transformed in various ways through a diverse collection of program modules (models). It is therefore necessary to have not only a comprehensive collection of such models, but also suitable mechanisms to use and control these models to solve problems that are usually phrased as queries to the system. In other words, the design of the model-management subsystem has a major impact on the interaction between a user and the DSS. This article addresses two key issues in model-management system design: machine representation of models and development of mechanical methods for their manipulation.

Engelman, C., C. H. Berg, and M. Bischoff, "KNOBS: An Experimental Knowledge Based Tactical Air Mission Planning System," Proceedings of the Sixth IJCAI, American Association for Artificial Intelligence, 1979.

KNOBS (from KNOwledge Based System) refers to an experimental system under development at MITRE corporation to explore the applicability of artificial intelligence techniques to the implementation of an automated, extremely flexible tactical mission planning consultant. This paper discusses that system and another one which supports the rule-based simulation of aircraft identification strategies.

Ercegovac, Z., "Knowledge-Based Expert Systems: A Profile and Implications," National Online Meeting Proceedings, Apr. 1984.

The purpose of this article is to discuss some of the implications of knowledge-based expert systems in general, and specifically the applications in library-related practice. The main foci of the article are: general profile of research and development in knowledge-based expert systems; implication of knowledge-based expert systems in the context of library practice, and the systems' role for the future management of information resources, services, and products; and impact of artificial intelligence on professional goals and values for excellence in library services, particularly in view of the financial pressures that face libraries today. It is anticipated that this article will be useful library services, particularly in view of the financial pressures that face libraries today. It is anticipated that this article will be useful to libraries and users of library services who will be affected by the AI technology as it unfolds.

Fallon, R., "Rule-Based Modeling as an Analysis Tool: Implications for Resource Allocation Within the Strategic Air Command," Report No.: N-1489-AF, RAND Corporation, Santa Monica, CA, 1980.

Fieschi, M., "Expert Systems: A New Look at Computers in Medical Education, Meeting the Challenge: Informatics and Medical Education," Proceedings of the IFIP-IMIA Working Conference, Amsterdam, Netherlands: North-Holland, 1983.

The classic medical education aid programs are directed mainly towards teaching objectives aimed at the acquisition of knowledge; they are rigid and inflexible where the acquisition of ability is concerned. Artificial intelligence methods and expert systems make it possible to contemplate the construction of systems which have knowledge of the field being studied, of the pedagogical data, of the student's level, and knowledge allowing a man-machine communication very close to natural language. Moreover, by their construction, these systems present considerable advantages with regard to the modularity and easy up-dating of the knowledge bases. These facilities, in addition to the high level of expertise they present on the subject dealt with, allow one to try different explicit teaching strategies and to provide explanations adapted to the student.

Forbus, K. D., "Qualitative Reasoning about Physical Processes," Proceedings of the Seventh IJCAI, American Association for Artificial Intelligence, 1981.

Fox, J., and A. Rector, "Expert Systems for Primary Medical Care?" Automedica, Vol. 4, No. 2/3, 1982.

Fujita, T., and S. Goto, "A Rule-based Routing System," Proceedings of the IEEE International Conference on Computer Design: VLSI in Computers, New York: IEEE, 1983.

Gale, W. A., and D. Pregibon, "Artificial Intelligence Research in Statistics," AT&T Bell Laboratories, Murray Hill, NJ, 1984.

Gale, W. A., and D. Pregibon, "An Expert System for Regression Analysis," Proceedings of the Fourteenth Symposium on the Interface of Computer Science and Statistics, New York, NY: Springer-Verlag, 1982.

Gaschnig, J., "Prospector: An Expert System for Mineral Exploration," in D. Michie (Ed.), Introductory Readings in Expert Systems, New York, NY: Gordon Breach Science Publishers, 1982.

Gaschnig, J., "Application of the Prospector System to Geological Exploration Problems" in J. E. Hayes, D. Michie, and Y. H. Pao (Eds.), Machine Intelligence 10, Chichester, England: Horwood, 1982.

A partial criterion for the success of a knowledge-based problem-solving system is its usefulness as a tool to those working in its specialized domain of expertise. Here we describe several applications of the PROSPECTOR consultation system to mineral exploration tasks. One was a pilot study conducted for the National Uranium Resource Estimate program of U.S. Department of Energy. This application estimated the favorability of several test regions for occurrence of sandstone uranium deposits. For credibility, the study was preceded by a performance evaluation of the relevant portion of PROSPECTOR's knowledge base, which showed that PROSPECTOR's conclusions agreed very closely with those of the model designer over a broad range of conditions and levels of detail.

Gelernter, H. L., et al., "Empirical Explorations of SYNCHEM," Science, Vol. 197, 1977.

Gershman, A., "Building a Geological Expert System for Dipmeter Interpretation," Proceedings of the ECAI-82, 1982.

Geschke, M. J., R. A. Bullock, and L. E. Widmaier, "TAC II: An Expert Knowledge-Based System for Tactical Decision Making," Master's thesis, Naval Postgraduate School, Monterey, CA, June 1983.

There exists a genuine need for a tactical decision-making system within the Department of Defense for the small scale environment tactical decision maker. To this end, we propose TAC II, a prototypical system for tactical decision making, to be implemented as a distributed system on microcomputers. TAC II is a redesign and partial implementation of an expert Artificial Intelligence system proposed by previous Naval Postgraduate School students. The system receives preprocessed sensor inputs, determines what contacts are present, and suggests the best possible actions to take. It performs target analysis and correlation based on the current tactical situation. Production rules are used to discover which actions have been established by higher authority for the current tactical situation. A pattern matching algorithm provides a heuristic means of identifying similar known situations, and suggests actions to take based on those situations.

Goldstein, I. P., "The Genetic Graph: A Representation for the Evolution of Procedural Knowledge," in D. Sleeman and J. S. Brown (Eds.), Intelligent Tutoring Systems, New York: Academic Press, 1982.

Gorry, G. A., H. Silverman, and S. G. Pauker, "Capturing Clinical Expertise: A Computer Program that Considers Clinical Responses to Digitalis," American Journal of Medicine, Vol. 64, 1978.

Groundwater, E. H., "A Demonstration of an Ocean Surveillance Information Fusion Expert System," _Proceedings of the SPIE International Society of Optical Engineering_, Vol. 485, May 1984.

The expert system described here models the thought processes of an ocean surveillance watch analyst attempting to assess vessels' missions and destinations, given their correlated tracks, history, and location/status of other vessels in the domain of interest. The expert system has been developed using the OPS5 production system and the Franz LISP programming language on a VAX/VMS computer system. The architecture and design of the expert system are presented, implementation issues are discussed, and progress and future plans are reviewed.

Hahn, G. J., "More Intelligent Statistical Software and Statistical Software and Statistical Expert Systems: Future Directions," _The American Statistician_, Vol. 39, No. 1, Feb. 1985.

Statistical computer programs are becoming increasingly accessible to people with limited statistical training. More intelligent statistical software is clearly needed. In this article, new or improved offerings -- ranging from computer-based indexes of the literature to expert statistical systems -- are discussed and illustrated. Three general levels of statistical software are differentiated: computerized statistical answering and referral services, expert guidance embedded in statistical programs, and automated statistical consultation and data analysis. The first of these is already feasible and is being implemented in various ways. Embedding expert guidance in statistical programs is a technically challenging, but highly worthwhile undertaking. Automating statistical consulting and data analysis in the form of a statistical expert system is fraught with difficulties; the best chances of success appear to be in researching specialized modules. Possible features of an expert system module for product-life data analysis are briefly described.

Harandi, M. T., "A Knowledge-Based Programming Support Tool," _Proceedings of Trends and Applications 1983 on Automating Intelligent Behavior: Applications and Frontiers_, Silver Spring, MD: IEEE Computer Society Press, May 1983.

This paper presents an overview of a knowledge-based programming support tool. Although the system would not synthesize programs automatically, it has the capability of aiding programmers in various phases of program production such as design, coding, debugging and testing. The underlying design principles of this system are similar to those governing the implementation of knowledge-based expertise in other domains of human mental skill. The system is composed of several major units, each an expert system for a sub-domain of program development process. It implements various elements of programming expertise by an interactive system which allows the domain specialist to easily and effectively transfer to the system the knowledge it needs for its decision making.

Harandi, M. T., "Knowledge-Based Program Debugging: A Heuristic Model," _Proceedings of the Conference on Software Development Tools,_ Techniques, and Alternatives, Silver Spring, MD: IEEE Computer Society Press, July 1983.

Program debugging is a time and resource consuming task. The fact that there are a large number of human experts in this domain indicates that this is an area where knowledge-based techniques can be applied. This paper discusses a knowledge-based program debugging system which embodies heuristics dealing with a majority of compile time and certain run time errors for which sufficient explicit symptoms are present. The paper includes a brief description of expert systems, an overview of the properties of the domain of program debugging, and a presentation of the structure and function of the debugging system.

Hartley, R. T., "CRIB: Computer Fault-Finding Through Knowledge Engineering," Computer, Vol. 17, No. 3, Mar. 1984.

Hayes-Roth, F., and V. R. Lesser, "Focus of Attention in the Hearsay-II Speech Understanding System (Interim report)," Report No.: AFOSR-TR-77-0326, Dept. of Computer Science, Carnegie-Mellon Univ., Pittsburgh, PA, Jan. 1977.

Hearsay-II is a complex, distributed-logic system for speech understanding developed at Carnegie-Mellon University. Processing is performed by independent, data-directed knowledge source processes that examine and alter values in a global data base representing hypothesized phonetic segments, syllables, words, and phrases, as well as the hypothetical temporal and logical relationships between them. The question of how the numerous potential activities of the knowledge sources should be scheduled to complete the interpretation of an utterance in minimal time is called the focus of attention problem. Near optimal focusing is especially important in a speech understanding system because of the potentially very large solution space that needs to be searched. This focus of attention problem is representative of general resource allocation problems involving cooperative and competitive processes. A general attention control mechanism is developed that facilitates the experimental evaluation of a variety of specific attention control policies (such as best-first, bottom-up, and top-down search strategies) and allows the modular addition of specialized heuristics for the speech understanding task. Empirical results demonstrate the effectiveness of the focusing principles. Possible directions for future research are considered.

Hertz, D. B., "Artificial Intelligence and the Business Manager," Computerworld, Vol. 17, No. 43, Oct. 1983.

The Fortune 1000 executive may have too much data too soon. His problem lies in turning it into usable information. Artificial intelligence (AI) may promise one resolution of the executive's dilemma. Knowledge resides in the facts; mining that knowledge for more effective competition, smoother operations and better planning is a major goal of AI knowledge scientists and engineers. They believe the right kind of computer programs will turn the tide of data from an uncontrollable flood to an orderly, productive stream.

Hollander, C. R., et al., "The Drilling Advisor," Proceedings of Trends and Applications 1983 on Automating Intelligent Behavior: Applications and Frontiers, Silver Spring, MD: IEEE Computer Society Press, May 1983.

This paper discusses the Drilling Advisor, an interactive computer-based consultant whose goal is to provide advice on how to resolve problems encountered in the course of drilling production oil wells. At present, only problems which manifest themselves as striking of the drilling mechanism within the bore-hole are considered. The Drilling Advisor is an example of a knowledge-based expert system in which a superior level of task performance is achieved by representing, in data structures, a huge amount of expert knowledge obtained from a human specialist.

Horn, W., W. Buchstaller, and R. Trappl, "Knowledge Structure Definition for an Expert System in Primary Medical Care,"Proceedings of the Seventh IJCAI, American Association for Artificial Intelligence, 1981.

Hudson, D. L., and T. Estrin, "EMERGE - A Data-driven Medical Decision Making Aid," IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 6, No. 1, Jan. 1984.

EMERGE is an expert system designed as a medical decision-making aid. It is machine-dependent, and is implemented in standard PASCAL. It has modest memory requirements and can operate on a microcomputer. EMERGE is rule-based and its initial application is the analysis of chest pain in the emergency room. The knowledge base is maintained separately from the consultation program. Thus the application area can be changed without any modification to the software. This paper describes the control structures and rule searching procedures used in EMERGE.

Hughes, P., "Keyword Access System," Creative Computing, Vol. 6, No. 3, March 1980.

The Keyword Access System (KAS) is a set of programs designed to store information about magazine articles and allows this information to be retrieved in various ways. The data in the file used by KAS consist of article name, author, magazine name, date, page number, and up to ten keywords or descriptors. The purpose of KAS is to access information more easily either by an online inquiry using a terminal, or by looking through a printed report. The online inquiry program allows the user to qualify a search by any or all of the following: magazine name, author, up to 10 keywords. A printed report of all the data is available in the order the information is stored in the file and, also, sorted by keyword. KAS is implemented on a SWTPC MP-68 computer using Basic 3.0.

Jones, M., "Applications of Artificial Intelligence Within Education," Computers and Mathematics, Volume 11, No. 5, 1985.

Jones, M. J., and D. T. Crates, "Expert Systems and Videotex: An Application in the Marketing of Agrochemicals," Proceedings of Videotex '84 International, London, England: Online Publications, 1984.

The development of a computer-based agricultural advisory system is described bringing together the technologies of expert systems and videotex. The factors leading to the choice of these media are discussed in the context of the role of the system as a marketing aid.

Johnson, W. L., and E. Soloway, "PROUST: Knowledge-Based Program Understanding," Report No.: YALEU/DCS/RR-285, Department of Computer Science, Yale University, New Haven, CT, Aug. 1983.

This paper describes a program called PROUST which does on-line analysis of PASCAL programs written by novice programmers. PROUST takes as input a program and a nonalgorithm description of the program requirements, and finds the most likely mapping between the requirements and the code. This mapping is in essence a reconstruction of the design and implementation steps that the programmer went through in writing the program. A knowledge base of programming plans and strategies, together with common bugs associated with them, is used in constructing this mapping. Bugs are discovered in the process of relating plans to the code. PROUST can give explanations of program bugs by relating the buggy code to its underlying intentions.

Kaihara, S., et al., "A Rule-Based Physicians' Consultation System for Cardiovascular Diseases," Proceedings of the International Conference on Cybernetics and Society, New York: IEEE, Oct. 1978.

Kohler, J. L., "The VENTILATION EXPERT for Underground Coal Mines," Proceedings of the First Symposium on the Theory and Application of Expert Systems in Emergency Management Operations, Special Publication 717, Gaithersburg, MD: National Bureau of Standards, 1986.

Kulikowski, C. A., "Artificial Intelligence Methods and Systems for Medical Consultation," IEEE Transactions on Pattern Analysis Machine Intelligence, Vol. 2, No. 5, Sept. 1980.

The major artificial intelligence (AI) problems that arise in designing a consultation program involve choices of knowledge representations, diagnostic interpretation strategies, and treatment planning strategies. The need to justify decisions and update the knowledge base in the light of new research findings places a premium on the modularity of a representation and the ease with which its reasoning procedures can be explained. An important insight that has resulted from the design of several artificial intelligence systems is that robustness of performance, in the presence of many uncertainty relationships, can be achieved by eliciting from the expert a segmentation of knowledge that also provides a rich network of deterministic relationships to interweave in the space of hypotheses. A number of knowledge-based AI representational schemes that generalize the results of the early consultation programs have emerged recently (EMYCIN, EXPERT, AGE). They provide techniques and tools for system building that promise to be very versatile in helping to design new consultation programs.

Kulikowski, C. A., and S. M. Weiss, "Representation of Expert Knowledge for Consultant," in P. Szolovits (Ed.), Artificial Intelligence in Medicine, Boulder, CO: Westview Press, 1982.

Lavrac, N., "KARDIO-E - An Expert System for Electrocardiographic Diagnosis of Cardiac Arrhythmias," Expert Systems, Vol. 2, No. 1, Jan. 1985.

KARDIO-E is an expert system for the electrocardiographic diagnosis of heart arrhythmias. At the present stage, the system can be used as a helpful diagnostic tool in the routine assessment of ECG recordings in preventive or systematic examinations as, by a cardiologist's estimation, its performance is equivalent to that of a specialist of internal medicine (non-cardiologist) highly skilled in the reading of ECG recordings. The system is also useful for instruction in electrocardiography education. This paper presents the ECG diagnosis problem, the automatic generation of the system's knowledge base, and a detailed description of the system. Results of an assessment study are also presented.

Lenat, D. B., "EURISKO: A Program that Learns New Heuristics and Domain Concepts," Artificial Intelligence, Volume 21, 1983.

Lewis, J. W., and F. S. Lynch, "GETREE: A Knowledge Management Tool," Proceedings of Trends and Applications 1983 on Automating Intelligent Behavior: Applications and Frontiers, Silver Spring, MD: IEEE Computer Society Press, May 1983.

In practical expert systems, regardless of inference technology, the accuracy and completeness of the knowledge base determine expert system performance, and the cost of acquiring that knowledge base tends to dominate all other hardware and software costs. To reduce knowledge acquisition cost and error rate, GETREE -- a new interactive knowledge management system -- is being designed and implemented in GE Corporate Research. In the new system, the knowledge base is represented as an (almost tree-like) network of nodes which can be displayed and manipulated on a personal computer workstation. Users -- even unsophisticated users -- can build and navigate these networks, modify nodes and connecting arcs, verify correctness visually, follow the execution of inference engines, and generate equivalent code to be run in more constrained target environments. Consequently, with GETREE, the user-organization can have full responsibility for and control of the knowledge base.

Lumley, J. W., "Using Expert Systems for Planning," IEE Colloquium on Decision Support Aspects of Expert Systems (Digest No. 67), London, England: IEE, May 1984.

Planning is an important part of many decision-making processes involving generation of a sequence of actions, prediction of the effect of these actions, and modifications of the plan in the light of these predictions. As expert systems begin to be used for more demanding decisions, the ability to plan becomes essential. This article considers the problems involved with using expert systems in this planning role.

McDermott, J., "XSEL: A Computer Salesperson's Assistant," in J. E. Hayes, D. Michie, and Y. H. Pao (Eds.), Machine Intelligence 10, Chichester, England: Horwood, 1982.

R1, a knowledge-based configuere of VAX-11 computer systems, has been used by Digital Equipment Corporation's Manufacturing Organization. The success of this program, and the existence at DEC of a newly formed group capable of supporting knowledge-based programs, has led other groups at DEC to support the development of programs that can be used in conjunction with R1. This paper describes XSEL, a program being developed at Carnegie-Mellon University that will assist salespeople in tailoring computer systems to fit the needs of customers. XSEL will have two kinds of expertise: it will know how to select hardware and software components that will fulfill the requirements of a particular set of applications, and it will know how to provide explanations in the computer system sales domain.

McKeown, Jr., M. D., and J. McDermott, "Toward Expert Systems for Photo Interpretation," Proceedings of Trends and Applications 1983 on Automating Intelligent Behavior: Applications and Frontiers, Silver Spring, MD: IEEE Computer Society Press, May 1983.

This paper describes some preliminary results in the design and implementation of a system for semi-automatic photo-interpretation of high resolution aerial photographs. The system, SPAM, consists of three major components, an image/map database, a collection of image processing tools, and a rule-based system whose domain of expertise is commercial airports in general, and National Airport (Washington, D.C.) in particular. Discussed are the design rationale and implementation. Applications for such photo-interpretation systems include cartography and decision-support systems for situation assessment.

Michalski, R. S., and R. L. Chilausky, "Learning by Being Told and Learning from Examples: An Experimental Comparison of Two Methods of Knowledge Acquisition in the Context of Developing an Expert System for Soybean Disease Diagnosis," International Journal of Policy Analysis and Information Systems, Volume 4, No. 2, Dec. 1980.

Mill, J., "Expert Systems Made Simple," Computer Magazine, Feb. 1985.

Miller, F. D., et al., "ACE (Automated Cable Expert) Experiment: Initial Evaluation of an Expert System for Preventive Maintenance," Proceedings of the Joint Services Workshop on Artificial Intelligence in Maintenance, Boulder, CO, Oct. 1983.

ACE is a knowledge-based system designed to direct preventive maintenance activities in the local telephone network. ACE was developed at AT&T Bell Laboratories to demonstrate and evaluate the potential for using expert systems technology in cable maintenance. This paper describes the role of ACE in the maintenance process and discusses the methods used to evaluate ACE's performance in the field. Some general issues for the evaluation of expert systems are also considered.

Mittal, S., B. Chandrasekaran, and J. Sticklen, "PATREC: A Knowledge-Directed Database for a Diagnostic Expert System," Computer, Vol. 17, No. 9, Sept. 1984.

PATREC is one component in a cluster of knowledge-based medical expert systems developed at Ohio State University. This system holds extensive knowledge on relevant medical data entities; it uses this knowledge to acquire and organize data about patients and to answer questions about patient data needed for diagnostic reasoning. The most important feature of Patrec is its use of inferential knowledge, embedded in the underlying knowledge base, to generate answers when corresponding data are not explicitly stored in the database. In addition, PATREC can handle some temporal aspects of the data and answer questions regarding the temporal relationships of events. While Patrec is an experimental vehicle using a variety of state-of-the-art AI representational techniques, it is, in fact, extensively used by an automated diagnosis system called MDX.

Negoita, C. V., "Management Applications of Expert Systems," Human Systems Management, Vol. 4, No. 4, Autumn 1984.

Only recently have expert systems advanced to the point that management applications are promising to accomplish practical results. Most of these results can be attributed to the fact that they can have natural language front ends and accept imprecise data. Internalizing a model of the external milieu, they avoid interfaces and permit a real dialogue with the users.

Nelson, D. E., "Combat Battle Damage Assessor Expert System," Report No.: SBI-AD-F650 094, Air Force Wright Aeronautical Labs., Wright-Patterson AFB, OH, May 1984.

Newell, A., et al., Speech Understanding Systems: Final Report of a Study Group, New York: American Elsevier, 1973.

Newell, A., "HARPY: Production Systems and Human Cognition," Report No. CMU-CS-78-140, Computer Science Department, Carnegie-Mellon University, Pittsburgh, PA, 1978.

Nii, H. P., and E. A. Feigenbaum, "Rule-Based Understanding of Signals," in D. Waterman and F. Hayes-Roth (Eds.), Pattern-Directed Inference Systems, New York: Academic Press, 1978.

Nii, H. P., et al., "Signal-to-Symbol Transformation: HASP/SIAP Case Study," AI Magazine, Volume 3, No. 2, 1982.

O'Keefe, R. M., "Expert Systems and Operational Research - Mutual Benefits," Journal of Operational Research Society, Vol. 36, No. 2, Feb. 1985.

The aim of much operational research and expert systems work is similar (i.e., helping a decision maker do his job). Liaison will produce mutual benefits. Expert systems development will benefit from the experience of operational research. The operational research community is well placed to support and develop the use of such systems by management. In this paper the author considers: knowledge as a model, expert systems in the OR domain, knowledge acquisition, and decision support systems.

O'Keefe, R. M., "An Expert System for Statistics," Technical Conference on Theory and Practice of Knowledge Based Systems, England: Brunel University, 1982.

Oleson, C. E., "EXAMINER: A System Using Contextual Knowledge for Analysis of Diagnostic Behavior," Proceedings of the Fifth IJCAI, American Association for Artificial Intelligence, 1977.

Pauker, S. et al., "Toward the Simulation of Clinical Cognition: Taking the Present Illness by Computer," American Journal of Medicine, Volume 60, 1976.

Pople, H. E., Jr., "Heuristic Methods for Imposing Structure on Ill-Structured Problems: The Structuring of Medical Diagnosis," in P. Szolovits (Ed.), Artificial Intelligence in Medicine, Boulder, CO: Westview Press, 1982.

Portier, K. M., and P. Y. Lai, "A Statistical Expert System for Analysis Determination," Proceedings of the Statistical Computing Section, American Statistical Association, 1983.

Reboh, R., "Knowledge Engineering Techniques and Tools in the Prospector Environment," SRI Technical Note 243, AI Center, SRI International, Menlo Park, CA, May 1982.

Reggia, et al., "Towards an Intelligent Textbook of Neurology," Proceedings of the Fourth Annual Symposium on Computer Applications in Medical Care, 1980.

Rosenberg, S., "Knowledge-Based System for Providing Intelligent Access to a Petroleum User Database," Dept. of Energy, Lawrence Berkeley Laboratory, University of California at Berkeley, 1979.

Rychener, M. D., "Expert Systems for Engineering Design," Expert Systems, Vol. 2, No. 1, Jan. 1985.

An expert system embodies a human expert's domain-specific knowledge and skill, acquired and refined over years of experience. A number of problems in engineering design can be solved by using current expert system techniques. This paper enumerates the main components of such problems and the steps that are taken in solving them. A few prototypical artificial intelligence systems embody techniques that can be applied to engineering problems. These are surveyed, and their relevance to components of design problems is discussed. Some expert systems in design domains are summarized, with emphasis on aspects that can illustrate wider applicability of the techniques. The area of engineering design offers rich opportunities for advancing the state-of-the-art in expert systems. An annotated bibliography is included.

Scarl, E. A., J. Jamieson, and C. Delaune, "Knowledge-Based Fault Monitoring and Diagnosis in Space Shuttle Propellant Loading," Proceedings of the 1984 IEEE National Aerospace and Electronics Conference (NAECON), Vol. 2, New york: IEEE, 1984.

LES (LOX Expert System) is a knowledge-based system being applied to a semi-real-time system application, monitoring the loading of a cryogenic fuel for the space shuttle. The system's design and fault isolation techniques are described.

Schutzer, D., "Applications of Artificial Intelligence to Military Communications," Proceedings of the 1983 IEEE Military Communications Conference, Vol. 3, New York: IEEE, 1983.

This paper explores the field of artificial intelligence with respect to its application to military communication design problems. In particular, it is shown how natural processing languages and knowledge-based system technologies can be used to reduce the required capacity, and to improve a communication system's robustness and tolerance of errors by trading off computation for communication. These technologies are also shown to improve the security of a military communications system. Other applications of artificial intelligence technology include the use of expert systems in the operation, control and maintenance, and training areas.

Shannon, T. C., "XCON - DEC's Artificial Intelligence VAX Configurer," VAX/RSTS Prof., Vol. 6, No. 5, Oct. 1984.

Shaw, R., "Artificial Intelligence and Personnel Planning," Proceedings of the Third National Conference and Exhibition on Computers in Personnel: Making Manpower Profitable, London, England: Institute of Personnel Management, June 1984.

Shortliffe, E. H., "MYCIN: A Knowledge-Based Computer Program Applied to Infectious Diseases," Proceedings of the First Annual Symposium on Computer Application in Medical Care, New York: IEEE, Oct. 1977.

A rule-based expert system is described which uses artificial intelligence techniques and a model of the interaction between physicians and human consultants to attempt to satisfy the demands of a user community that is often reluctant to experiment with computer technology. Experience to data has demonstrated that the program is efficient, relatively easy to use, and reliable in the domain of bacteria therapy selection. Future work will involve broadening and evaluating the program's expertise in other areas of infectious disease therapy. To that end, rules regarding diagnosis and treatment of meningitis have been written and are under evaluation.

Shortliffe, E. H., et al. "ONCOCIN: An Expert System for Oncology Protocol Management," Proceedings of the Seventh IJCAI, American Association for Artificial Intelligence, 1981.

Shortliffe, E. H., Computer-Based Medical Consultations: MYCIN, New York: American Elsevier, 1976.

Shortliffe, E. H., B. G. Buchanan, and E. A. Feigenbaum, "Knowledge Engineering for Medical Decision Making: A Review of Computer-Based Clinical Decision Aids," Proceedings of the IEEE, Vol. 67, No. 9, New York: IEEE, Sept. 1979.

Shubin, H., and J. W. Ulrich, "IDT: An Intelligent Diagnostic Tool," Proceedings of the National Conference on Artificial Intelligence, AAAI-82, American Association for Artificial Intelligence, 1982.

Sumner, G. C., "Knowledge-Based Systems Maintenance Applications (ATE)," Proceedings of the Autotestcon '82: the 1982 IEEE International Automatic Testing Conference, New York: IEEE, Oct. 1982.

Taylor, J. H., and D. K. Frederick, "An Expert System Architecture for Computer-Aided Control Engineering," Proceedings of the IEEE, Vol. 72, No. 12, New York: IEEE, Dec. 1984.

Expert or knowledge-based systems are defined as software environments designed to aid in solving problems that require high levels of expertise, 'reasoning' (inference), and heuristics. The authors propose the development of a rule-based expert system to create a third-generation man/machine environment for computer-aided control engineering (CACE). The focus is primarily on the high-level requirements for an improved CACE environment, and on the expert system concepts and structures that have been conceived to fulfill these needs. The chief goal is to determine what artificial intelligence has to contribute to such an environment, and to provide as definite and as credible a vision of an expert system for CACE as possible. The main product of this effort is an expert system architecture for CACE. The authors close with a brief status report.

Thomas, R. C., "Knowledge-Based Techniques for Decision Support for Management," IEE Colloquium on Decision Support Aspects of Expert Systems (Digest No. 67), London: IEE, May 1984.

Tou, J. T., and J. M. Cheng, "Design of a Knowledge-Based Expert System for Applications in Agriculture," Proceedings of Trends and Applications 1983 on Automating Intelligent Behavior: Applications and Frontiers, Silver Spring, MD: IEEE Computer Society Press, May 1983.

This paper presents the design of a computer-based knowledge system, APRIKS, for application in agriculture. The APRIKS system is a pilot system which demonstrates the use of a minicomputer in agricultural knowledge transfer and utilization for diagnostic consultation, crop pest control, plant disease treatment, and management recommendations. The major components of this expert system consist of a knowledge base and a recognition/inference scheme. The design makes use of pattern-directed approach, as well as rule-based approach. The system can respond to both menu-selection inputs and simple natural language inputs. It performs three modes of operations: interactive browsing, diagnostic consultation, and question-answering.

Tsukiyama, M., and T. Fukuda, "An Application of Knowledge Base to Control Systems," *Proceedings of the International Conference on Cybernetics and Society*, New York: IEEE, Oct. 1981.

This paper discusses the knowledge base of plant operations and the design of knowledge base control system (KBCS) to realize an expert system for plant control. The structure of the KBCS is represented in a network organization of modules which comprise knowledge and calculation tools. The total structure is constructed by a set of rules describing the relation between two modules. An illustrative example is shown.

Underwood, W. E., "A CSA Model-Based Nuclear Power Plant Consultant," *Proceedings of the National Conference on Artificial Intelligence, AAAI-82*, American Association for Artificial Intelligence, 1982.

van Melle, W., "MYCIN: A Knowledge-Based Consultation Program for Infectious Disease Diagnosis," *International Journal of Man-Machine Studies*, Vol. 10, 1978.

MYCIN is a computer-based consultation system designed to assist physicians in the diagnosis of and therapy selection for patients with bacterial infections. In addition to the consultation system itself, MYCIN contains an explanation system which can answer simple English questions in order to justify its advice or educate the user. The system's knowledge is encoded in the form of some 350 production rules which embody the clinical decision criteria of infectious disease experts. Much of MYCIN's power derives from the modular, highly stylized nature of these decision rules, enabling the system to dissect its own reasoning and allowing easy modification of the knowledge base.

van Melle, W., "A Domain-Independent Production-Rule System for Consultation Programs," *Proceedings of the Sixth IJCAI*, American Association of Artificial Intelligence, 1979.

Vesonder, G. T., et al., "ACE: An Expert System for Telephone Cable Maintenance," *Proceedings of the Eighth IJCAI*, American Association for Artificial Intelligence, 1983.

ACE, a system for Automated Cable Expertise, is a knowledge-based expert system designed to provide trouble-shooting reports and management analyses for telephone cable maintenance. Design decisions faced during the construction of ACE were guided by recent successes in expert systems technology, most notably R1/XCON, the Digital Equipment Corporation VAX configuration program. ACE departs from standard expert system architectures in its use of a conventional data base management system as its primary source of information. Its primary sources of knowledge are the users of the database system and primers on maintenance analysis strategies.

Weiss, S., C. A. Kulikowski, and A. Safir, "A Model-Based Consultation System for the Long-Term Management of Glaucoma," Proceedings of the Fifth IJCAI, American Association for Artificial Intelligence, 1977.

Weiss, S., et al., "Building Expert Systems for Controlling Complex Programs," Proceedings of the National Conference on Artificial Intelligence, AAAI-82, American Association for Artificial Intelligence, 1982.

Wesson, R. B., "Planning in the World of the Air Traffic Controller," Proceedings of the Fifth IJCAI, American Association for Artificial Intelligence, 1977.

An enroute air traffic control (ATC) simulation has provided the basis for research into the marriage of discrete simulation and artificial intelligence techniques. A program which simulates, using real world data, the movement of aircraft in an ATC environment forms a robot's world model. Using a production system to respond to events in the simulated world, the robot is able to look ahead and form a plan of instructions which guarantees safe, expedient aircraft transit. A distinction is made between the real world, where pilots can make mistakes, change their minds, etc., and an idealized plan-ahead world which the robot uses. The overall simulation alternates between updating the real world and planning in the idealized one to investigate the robot's ability to plan in the face of uncertainty.

# 9. USE OF EXPERT SYSTEMS IN THE DOMAIN OF EMERGENCY MANAGEMENT

## 9.1 THE SUITABILITY OF EXPERT SYSTEMS FOR EMERGENCY MANAGEMENT PROBLEMS

Researchers have recognized that many of the problems of emergency management can be viewed as data management decision problems like those found in other applications areas. Thus, we have computer-based decision systems for emergency management utilizing the standard information and decision support procedures as tools for disaster policy analysis (Kunreuther et al. [1978]) and disaster management and support (Wallace and de Balogh [1985]; Belardo and Karwan [1984]; Belardo et al. [1982]). With the advent and popularization of expert systems, many of the decision requirements of emergency management were found to fit the philosophy and capabilities of current expert systems technology. In particular, expert systems have been developed for hazardous materials spills (Allen et al. [1982]; Johnson and Jordan [1983]), coal mine ventilation (Kohler [1986]), reactor accident (Nelson [1982, 1984]) and earthquake early-warning and assessment (de Balogh et al. [1983], Russo and Wilson [1983]). Although the applications and the end uses of the associated expert systems are few in number, these attempts testify to the suitability of expert systems to address a variety of emergency management problems. In general, however, it is felt that expert systems can be of most value in the response phase of emergency situations. We expand upon this view in the next section.

## 9.2 EXPERT SYSTEMS FOR RESPONSE

In Figure 9.1, we list a number of emergency problems and the expert system area of application for the response phase of each problem. It is in this phase that we feel expert systems have the greatest potential. Most of the response-oriented expert systems would be operated by State and local agencies, while at the national level, we would have FEMA expert systems for mobilization and enemy attack and recovery.

With reference to Figure 9.1, response expert systems appear to be applicable to the following areas:

1. The problems of moving people, material and equipment into and out of an emergency zone (evacuation routing)

2. The location and assignment of evacuees to emergency shelters (shelter allocation)

3. The determination that an emergency situation is likely to occur (early warning)

4. The systemic and decision processes by which a determination is made that an emergency is occurring and automatic controls are to be initiated to shut down or degrade essential equipment and services (automated response/shutdown)

## Figure 9.1 Emergency/Expert Systems Areas of Application

EXPERT SYSTEM AREA OF APPLICATION

| EMERGENCY | Evacuation routing | Shelter allocation | Early warning | Automated response/ shut down | Resource management (Presonnel, equipment, communications, information processing) | Fault (source) analysis |
|---|---|---|---|---|---|---|
| Nuclear plant disaster | ● | ● | | ● | ● | ● |
| Earthquake | ● | ● | | ● | ● | |
| Hurricane | ● | ● | ● | | ● | |
| Floods | ● | ● | ● | | ● | |
| Volcanoes | | | ● | | ● | |
| Enemy attack | ● | ● | | | ● | |
| Civil disorder | | | | | ● | |
| Terrorist attack | | | | | ● | |
| Epidemic | | | | | ● | ● |
| Power failure | | | | | ● | ● |
| Resource shortage | | | | | ● | |
| Economic emergency | | | | | ● | |
| Agricultural emergency | | | | | ● | |
| Boating, ships, ferry accidents | | | | | ● | |
| Airplane crash | | | | | ● | ● |
| Fire | ● | ● | ● | | ● | ● |
| Explosion | ● | | | | ● | ● |
| Mine disaster | | | | ● | ● | ● |
| Hazardous spills | ● | ● | | | ● | ● |
| Toxic airborne release | ● | ● | ● | | ● | ● |
| Transportation disaster | | | | | ● | ● |
| Dam failure | ● | ● | ● | | ● | ● |
| Water shortage/pollution | | | | | ● | ● |
| Air pollution | ● | ● | ● | | ● | ● |

5. The problems associated with marshaling of the full range of resources necessary for the proper response to an emergency (resource management)

6. The analysis and determination of the source, type and extent of certain emergency situations (fault analysis)

We next discuss items (1), (2) and (5) in terms of possible activities and research in expert systems. As noted below, expert system requirements for these areas are interrelated with shelter allocation dependent on routing decisions, and resource management dependent on both shelter allocation and routing decisions.

The problem of evacuation routing is critical in those emergencies that involve the general public to a high degree. Recent chemical spills and toxic airborne releases have captured the headlines and have raised serious questions as to whether communities that live in the shadow of chemical plants are in a position to respond to a serious event. Thus, there is a need to establish and test contingency plans for notifying the surrounding community that an emergency has occurred, with the initiation of specific plans based on an expert system's analysis of the emergency (e.g., weather, time of day and other critical factors). The development of a prototype expert system for evacuation routing, designed and tested for a specific community, should be investigated. The extensive literature in routing procedures would be of value here.

For those emergencies in which evacuation is required, shelter allocation must be coordinated with routing assignments. Hence, a more general expert system that integrates the emergency management functions of escape routing and shelter allocation would be of value. The requirements of a prototype expert system that includes both routing and allocation should be studied for a high-risk community.

As evidenced by Figure 9.1, the problem of general resource management (personnel, equipment, communications, information processing) is one that cuts across all emergencies; it represents a fruitful area for investigation. As a generic expert system for resource management that is applicable to the full range of emergencies is not possible, attention should be focused on resource problems associated with evacuation type problems (e.g., toxic release, hurricane). For these problems, we feel that the expert system requirements of resource management and those of routing and shelter allocation can be combined into a powerful emergency control tool. Further, such an integrated system, when coupled with local and regional emergency information and data systems, would be able to assess changing conditions and thus, be an extremely valuable decision aid to local emergency managers.

In line with the above discussion, we cite the following bibliographic items: Belardo et al. [1984] and de Balogh [1986] for managing disasters using decision support systems and microcomputers; Petrie and Hathaway [1984] for automated response to high-speed rail accidents; Ferrell and Krzysztofowicz [1983] and Cuena [1983] for flood warning and dam control, and Heaton [1985] and Moore et al. [1985] for earthquake warning and assessment. Specifications for such an expert system based on the requirements of a specific community should be developed.

Within the other phases of emergency management -- mitigation, preparedness and recovery -- expert systems can also be of value. For example, expert systems can play an important role in preparedness if they can be integrated as decision aids into training exercises, with the expert systems being those that the trainee would actually use in the field. During a recovery phase, expert systems can be used to analyze eligibility for loans and government reconstruction programs.

## 9.3 MAJOR EXPERT SYSTEMS RESEARCH TOPICS

In this section, we describe expert system applications in areas that stem from FEMA's operational and educational activities; these areas are of broad scope and/or more technically oriented than the ones described above. In particular, we address how expert systems can be used by FEMA in the analysis of computer-based model outputs, automated early warning, training and research, and information retrieval; what areas of expert system research would be of most value to FEMA; and conclude with a short list of general expert system activities that should be considered by FEMA.

- Model Output Reviewer

  FEMA uses many large-scale econometric, input-output and other models to analyze impacts of emergencies on the nation's industrial and transportation system. Such analyses are often based on changing scenarios and require the review of a great deal of computer output. It has been proposed that FEMA investigate a computer-based process by which such outputs can be reviewed automatically, and the essence of the data be processed and formatted by an expert system that replicates what the human expert analysts would accomplish. At a minimum, such an expert system would be able to organize the model outputs into tables and graphs; at a maximum, it would be able to find, filter and adjust conflicting data and to make recommendations for future analyses and specific actions. Output analysis programs have been used for quite some time, but here we are proposing a major advance in such post-processor programs. Issues that would have to be resolved include the specification of the models, the associated expert(s) to be used to build a knowledge base, and the types of modeling exercises to be reviewed.

- Automated Early Warning/Response

  Some emergencies such as earthquakes or nuclear power-plant accidents require a means for early detection and, given that there is no means of stopping the incident, a quick response (e.g., within 100 seconds) is required to shutdown or softly degrade critical systems and to warn emergency managers. It has been suggested that an expert system, tied into a detection network, would be able to automatically shutdown critical parts of the infrastructure and issue the appropriate warnings. The expert system would be designed to interpret the information produced by the detection network and to make the "call" that a destructive incident is about to occur. There are a number of issues here that go beyond just being able to build such an expert system. For example, our present ability to differentiate earthquake signals is not perfected and false alarms would trigger needless actions (although the expert system triggering mechanism could be designed not to be activated unless an earthquake of sufficient Richter magnitude was actually happening). Furthermore, critical facilities may not consent to be tied into such an automatic system. From an expert system point of view, the logic design would need a proper research base and an ability to be tested and validated.

- Expert System for Training and Research

  It has been proposed to investigate how expert system concepts can be integrated within the present training program of the FEMA National Emergency Training Center (NETC). The current Central City Training exercise could be augmented by expert system decision aids for evacuation routing, allocation of resources, and shelter allocation. Using the NETC would enable FEMA to test out the utility of expert systems, and if successful, initiate development of operational systems for transfer to field locations. Research would be required to scope the application areas, and to design the decision rules and algorithmic processes for evaluation of the exercise information. This concept can also be extended to FEMA Headquarters' exercises in which expert systems can be used as decision aids and for automating certain exercise functions.

- Information/Knowledge Base Access

  Any emergency, especially at the national level, requires access to a variety of data bases and, given expert systems will be used in the emergency management area, access to many types of knowledge bases. The proper use of such information can be a complex matter and could possibly require an expert system, or certainly an advanced information retrieval system. At issue here is the structuring and availability of the data and knowledge bases, the usages and analyses to be made of the information, and the availability of the proper development tools.

- Related Research Topics

  As noted, the expert system/artificial intelligence (ES/AI) field has not reached its full potential and is under constant development. A number of ES/AI research topics that would hopefully be of benefit to the emergency management field have been proposed and are cited next.

-- Natural Language Dialogue Interface

Research in this area is needed to improve the ability to use natural English and the emergency management jargon for those who interface with computer and related information systems. The hope is to ease the use of computers by non-technical emergency management personnel by enabling them to hold a dialogue with the computer system to retrieve information and initiate analyses.

-- Decision Rules in Expert Systems

There are a number of basic areas of interest here such as:

(1) The manner in which an expert system analyzes uncertainty to determine certainty factors.

(2) The ability to state a full set of contingencies and decision rules.

(3) The way partial and imprecise information is interpreted.

(4) The means by which an expert system "learns" during its lifetime.

(5) The use of decision models, such as the Analytic Hierarchy Process, within an expert system.

-- Validation of Expert Systems

As with all decision models, there is a need to validate the recommendations of an expert system in terms of their consistency and acceptability by emergency managers, and their ability to replicate expert decisions. The process of validation is extremely important with respect to FEMA expert systems, as the automation of decisions that impact human lives is always open to second guessing and the clarity of hindsight. Thus, for an expert system, FEMA must establish a process with which it is validated by the associated experts and field-tested. The design of validation processes and the conducting of associated experiments are expert system aspects of prime concern to FEMA.

-- Expert Systems Shells and PC-Based Expert Systems

The building of an expert system can be facilitated greatly by using commercially available software and/or specialized equipment. Expert system shells (i.e., software that can handle the logic and construction of decision rules and is independent of the applications) need to be evaluated in terms of FEMA requirements. In addition, with the expected heavy use of personal computers (PC's) by field personnel, FEMA should evaluate commercial PC-based expert systems on field requirements, as well as FEMA Headquarters' applications.

## References

Allen Jr. J. D., et al., "Spills Problem and Applied Artificial Intelligence," Report No.: CONF-820418-23, Oak Ridge National Lab., Oak Ridge, TN, 1982.

This paper describes an Artificial Intelligence program designed for assisting in the location, assessment, and cleanup of hazardous spills and for the training of spills personnel who must deal with these matters. It includes a discussion of expert systems, a description of the system used, Carnegie-Mellon's OPS5, and illustrates a variety of rules, some taken directly from the SPILLS program.

Belardo, S., K. R. Karwan, and W. A. Wallace, "Managing the Response to Disasters Using Microcomputers," Interfaces, Vol. 14, No. 2, 1984.

The authors describe the microcomputer-based decision aids they developed that assisted four different disaster management organizations: the American Red Cross, the United States Coast Guard, a regional emergency medical organization, and the New York State Office of Disaster Preparedness. The four systems emphasize different components of a typical disaster response DSS.

Belardo, S., K. R. Karway, and W. A. Wallace, "An Investigation of System Design Considerations for Emergency Management Decision Support," IEEE Transactions on Systems, Management, and Cybernetics, Vol. SMC-14, No. 6, New York: IEEE, 1984.

Simulation tests carried out to evaluate DSS design factors in the context of emergency management decisionmaking are reported in this paper. In particular, the effects of two basic considerations are examined: 1) the availability of a microcomputer for enhancing data display during an emergency, and 2) the stage or severity of the emergency. The simulation was conducted at the training facility of the Federal Emergency Management Agency in Emmitsburg, MD with county civil preparedness personnel from around the country. The exercise involved the response to a nuclear generating facility accident, a situation that typically would require various inputs from public officials to protect the local community from potential radiation exposure. The results of the experiment indicate: 1) clear support for computerization of emergency management decision support systems, and 2) an apparent need to focus attention of further development of these systems to improve the decision-making capabilities of public managers during the later less-structured stages of an emergency.

Belardo, S., K. R. Karwan, and W. A. Wallace, "DSS Component Design Through Field Experimentation: An Application to Emergency Management," Proceedings of the Third International Conference on Information Systems, Ann Arbor, MI, 1982.

This paper describes the use of field experimentation in the development of a particular Emergency Management Decision Support System (EMDSS). The model component of this sytem, which contains a bi-objective transportation algorithm, was field tested with prospective users from a regional emergency medical organization. The results of the test are presented and implications for the use of field experimentation as a design strategy are discussed.

Berlin, G. N., "A Simulation Model for Assessing Building Firesafety," Model Systems, Inc., Boston, MA, undated.

This paper presents a computer simulation technique for analyzing emergency evacuation and building firesafety. The systematic, objective procedure includes mathematical modeling, graph theory and simulation to represent and analyze the building layout. Examples illustrating the application of this model to a mobile home are presented.

Brady, J., J. K. Bradford, and J. Mirolla, "Computer Applications for Emergency Management," Monograph No. 85-1, California Specialized Training Institute, San Luis Obispo, CA, 1985.

The purpose of this paper is to examine the use of computers in Emergency Management/Emergency Services. It analyzes the current and future use of computers in this field. The following questions are discussed: (1) Is the government using computers effectively? (2) Is the government spending a large amount of money on computers? (3) What kind of hardware and software is the government using?

Buzzell, C. A., et al., "Using Interactive Color Graphics to Model the Command Control Function in a Battlefield Simulation," in G. Birtwistle (Ed.), Artificial Intelligence, Graphics, and Simulation, La Jolla, CA: Simulation Councils, Inc., 1985.

The authors describe Janus, an event-driven, stochastic computer program simulating ground combat, in which players model the command/control function using interactive color graphics. Giving players the command/control function allows us to evaluate commanders' control measures and to understand their requirements for information. Players' actions and reactions are meant to simulate actual decision making on the battle field. As part of the simulation, the players may experiment with new methods; examine the effectiveness of their decisions; and study tactics, information flow, and weapons systems in various conflict scenarios. In the interactive mode, the simulation lacks repeatability. To offset this, players can war game to gain insight about a plan or weapon system and then run segments of the battle in a batch mode without interacting. This allows them to observe the sensitivity of the simulation to specific parameters.

Chartrand, R. L., "Information Technology for Emergency Management: The Many Potentials," CRS Review, 1984.

There is growing concern in many sectors of government and society that we as a nation -- at the federal level, as well as in the states and localities -- are inadequately prepared to cope with the technological and natural emergencies which affect people and property. The very complexity of our societal structure, as reflected in jurisdictional conflicts over responsiblity and the accelerated development and use of new energy sources, pesticides, and fertilizers with toxic side effects, have caused lawmakers and program managers to reexamine the effectiveness of our crisis (or emergency) management capability. One critical facet of improving this potentiality -- through improved mitigation, preparedness, response, and recovery mechanisms and procedures -- is the development of advanced information systems for acquiring, storing, processing, retrieving, and sharing essential data that may be used by emergency managers.

Chartrand, R. L., "Optimum Emergency Management: The Effective Use of Information Technology," Proceedings of the First Symposium on the Theory and Application of Expert Systems in Emergency Management Operations, Special Publication 717, Gaithersburg, MD: National Bureau of Standards, 1986.

Cuena, J., "The Use of Simulation Models and Human Advice to Build Expert System for the Defence and Control of River Floods," Proceedings of the Eighth IJCAI, American Association for Artificial Intelligence, 1983.

de Balogh, F., "Management Science Models and Disaster Management: Problems and Prospects for Capacity Building," Decision Analysis Series, Institute of Safety and Systems Management, University of Southern California, San Diego, CA, Mar. 1985.

This paper is designed to offer some observations as to the current status of model building in disaster management and the problems and prospects of such processes in enhancing the capacity of American society to deal with man-made and natural emergencies. Its perspective is based on the research being conducted by the Risk and Emergency Management Laboratory (REML) and Decision Support Systems Laboratory (DSSL) of the Institute of Safety and Systems Management of the University of Southern California (USC) in several types of disaster areas to include earthquakes, hazardous materials spills and discharges, and emergency evacuation. Specifically, this paper: (1) describes modeling and simulation approaches to two disaster types -- earthquake and hazardous materials spill/discharge -- currently being followed at USC, (2) uses these two examples of research to illustrate the problems associated with modeling these two types of disasters, and, (3) generalizes about the prospects of developing reliable and accurate models to enhance the disaster management capacity of governments. The paper concludes that, to date, management science models have not had a significant impact on the capacity of American society to deal with disasters. The problems of unreliable hazard risk estimation, complex and incomplete methodology, fragmented modeling research efforts, significant and costly data requirements, lack of transferability, imbalance in research resources committed to risk estimation models as opposed to models managing the response and training, and lack of a common program and plan of action to promote capacity building among public jurisdictions by federal research sponsoring agencies all combine to ensure that progress at best will continue to be atomized and incremental. It recommends some strategies for overcoming these problems.

de Balogh, F., "Decision Support and Expert Systems for Emergency Management Operations: A Microcomputer Approach," Proceedings of the First Symposium on the Theory and Application of Expert Systems in Emergency Management Operations, Special Publication 717, Gaithersburg, MD: National Bureau of Standards, 1986.

de Balogh, F., W. Petak, and J. Sessler, "A Computerized Demonstration Model for Earthquake Mitigation," Decision Support Systems Laboratory, University of Southern California, San Diego, CA, Feb. 1983.

This paper focuses on the potential power of utilizing computers for improved earthquake mitigation management within the framework of a new class of information system: decision support systems (DSS). A computerized demonstration system -- EDSS -- is described which is designed to extend the capability of planners and public officials to make effective decisions in earthquake mitigation. It consists of a micro-computer system which employs several data files on building structure inventories by planning areas and a damage calculation (modeling) component. The decision maker provides ground shaking intensity data for each planning area which is used by EDSS to calculate percentage of square feet damage and percentage of assessed value lost for each building classification and age group. The system permits "what if" analysis and is used interactively. A project of USC's Decision Support Sytems Laboratory (DSSL), EDSS is at present a demonstration model and not a comprehensive, working prototype DSS. Future plans call for expanding EDSS into a fully operational prototype available for testing by public officials and planners.

Ferrell, W. R., and R. Krzysztofowicz, "A Model of Human Response to Flood Warnings for System Evaluation," Water Resources Research, Vol. 19, No. 6, Dec. 1983.

A behavioral model of human response to flood warnings is developed as a component of a methodology for evaluation of the performance of flood forecast-response systems. A floodplain dweller responds to a sequence of flood warnings by taking protective action (such as evacuation, flood proofing, shutdown of a facility) in order to reduce his loss. The model is meant to mimic the actual response behavior of the floodplain dweller or to predict such behavior under future system conditions. It is building of mathematical representations of four interconnected cognitive elements of response: (1) uncertainty about flooding and loss prior to a flood, (2) sequential interference based on warnings during a flood, (3) response strategy, and (4) learning after a flood. Results of an application and qualitative verification of the model are discussed.

Gaffney, J. E., Jr. and I. R. Racer, "A Learning Interpretive Decision Algorithm for Severe Storm Forecasting Support," Proceedings of Trends and Applications 1983 on Automating Intelligent Behavior: Applications and Frontiers, Silver Spring, MD: IEEE Computer Society Press, May 1983.

As part of its ongoing program to develop new and better forecasting procedures and techniques, the National Weather Service has initiated an effort in "interpretive processing". Investigation has begun to determine the applicability of Artificial Intelligence (AI)/expert system algorithm that is being investigated to support the forecasting of severe thunderstorms.

Gaschnig, J., R. Reboh, and J. Reiter, "Development of a Knowledge-Based Expert System for Water Resource Problems," Technical Report SRI Project 1619, SRI International, Menlo Park, CA, Aug. 1981.

Gottingger, H. W., "HAZARD: An Expert System for Screening Environmental Chemicals on Carcinogenicity," Expert Systems, Vol. 1, No. 2, Oct. 1984.

The purpose of this paper is to design an expert system or an intelligent procedure to screen hazard potentials of environmental chemicals on the basis of structure-activity relationships in the study of chemical carcinogenesis. An analysis of a computerized database of known carcinogens (knowledge base) is being performed using the structure-activity trees in order to test the validity of the tree as a classification scheme (inference engine) and to evaluate trends or patterns that may exist between chemical structure and specificity for target tissue, route of administration, and animal species. Practical applications of the structure-activity tree depend on its eventual validation as a predictor of carcinogenic activity.

Heaton, T. H., "A Model for a Seismic Computerized Alert Network," Science, Vol. 228, May 1985.

In large earthquakes, damaging ground motions may occur at large epicentral distances. Because of the relatively slow speed of seismic waves, it is possible to construct a system to provide short-term warning (as much as several tens of seconds) of imminent strong ground motions from major earthquakes. Automated safety responses could be triggered by users after receiving estimates of the arrival time and strength of shaking expected at an individual site. Although warning times are likely to be short for areas greatly damaged by relatively numerous earthquakes of moderate size, large areas that experience very strong shaking during great earthquakes would receive longer warning times.

Hunt, A. R., and R. K. Mullen, "The Development of an Expert System to Assess the Vulnerability of Assets to Subnational Adversaries," Proceedings of the 1984 Carnahan Conference on Security Technology, New York: IEEE, May 1984.

Ishizuka, M., K. S. Fu, and J. T. P. Yao, "Rule-Based Damage Assessment System for Existing Sturctures," Report No.: CE-STR-82-14; NSF/CEE-82020, School of Civil Engineering, Purdue Univ., Lafayette, IN, Mar. 1982.

The state-of-the-art of damage assessment is reviewed and SPERIL version I, a rule-based damage system of existing structures, particularly those subjected to earthquake excitation, is described. The problem of damage assessment is considered in terms of the theory of pattern recognition and an expert system approach is used to develop the computer-based systems. In SPERIL version I, separate evidential observations are integrated on the basis of the extended Dempster and Shafer theory for subsets. Useful information is collected and expressed in a stylized rule format in the knowledge base. The current rules are expected to be updated by more specific ones for better performance.

Ishizuka, M., K. S. Fu, and J. T. P. Yao, "SPERIL I: Computer Based Structure Damage Assessment System," Report No.: CE-STR-81-36, Purdue University, Lafayette, IN, 1981.

Johnson, C. K., and S. R. Jordan, "Emergency Management of Inland Oil and Hazardous Chemical Spills: A Case Study in Knowledge Engineering," in F. Hayes-Roth, D. A. Waterman, and D. B. Lenat (Eds.), Building Expert Systems, Reading, MA: Addison-Wesley Publishing Company, 1983.

The authors of this article describe in some detail a typical problem in knowledge engineering: an expert system is needed to help consult with regular workers or to augment the limited experience of off-shift workers facing a difficult task. Like several other applications, this belongs to the class of crises-management problems. The oil and hazardous chemical spill-management task demands a system that can integrate diverse sources of knowledge, reason heuristically with incomplete and errorful data, accept data and advice continuously as they become available, and allocate limited resources to various tasks in a reasonable order.

Kohler, J. L., "The VENTILATION EXPERT for Underground Coal Mines," Proceedings of the First Symposium on the Theory and Application of Expert Systems in Emergency Management Operations, Special Publication 717, Gaithersburg, MD: National Bureau of Standards, 1986.

The underground coal mine environment does not present health or safety risks to workers when it is maintained to specific standards. Mining activity, as well as unforeseen events, can result in sudden and dangerous levels of explosive gas and dust. A mine may have hundreds of miles of passages spread over ten or more square miles. Control devices for maintaining the environment are distributed throughout the infrastructure of the mine. Many of the decisions which have to be made to maintain a safe environment are based on judgment, rather than documented facts or empirical procedures. All of these characteristics of the system compound to make it difficult to manage ventilation under routine conditions. When a fire or explosion occurs, the management problem escalates dramatically. The development of a prototype expert system to manage the operation of coal mine ventilation systems is described in this paper. The characteristics of the mine, as they affect the expert system design, are presented. The knowledge engineering aspects of this system, as well as implementation issues, are examined. The paper concludes with some observations on the requirements of a future version of the VENTILATION EXPERT.

Kunreuther, H., et al., "An Interactive Modeling System for Disaster Policy Analysis, Program on Technology, Environment and Man," Monograph #26, Institute of Behavioral Science, University of Colorado, 1978.

McColl, D. C., et al., "STAMMER2 Production System for Tactical Situation Assessment," Report No.: TD 298, Naval Ocean Systems Center, San Diego, CA, 1979.

Meitzler, W. D., and R. T. Jaske, "Exercise Evaluation and Simulation Facility," in G. Birtwistle (Ed.), Artificial Intelligence, Graphics, and Simulation, La Jolla, CA: Simulation Councils, Inc. 1985.

The authors describe the Integrated Emergency Management and Information System (IEMIS), a minicomputer based system that serves as a tool to assist FEMA in the evaluation of emergency plans and preparedness for natural and technological hazards. As an IEMIS subsystem, the Exercise Evaluation and Simulation Facility (EESF) provides capabilities speciifc to emergency exercises and plan development for application near commercial nuclear power facilities. EESF integrates the following resources: a meteorology model, evacuation model, dose model, maps, and exercise information into a single operational package. The user may assess these various resources, and on completion, display the results on a color graphic display or hardcopy unit. Key to the capability provided by EESF are the ease of use and quick comprehension of results presented via the graphic displays.

Mick, S., and W. A. Wallace, "Expert Systems as Decision Aids for Disaster Management," Report No.: 36-84-P6, Statistical, Management, and Information Sciences Department, Rensselaer Polytechnic Institute, Troy, NY, undated.

Moore, D., et al., "The FEMA Earthquake Damage and Loss Estimation System (FEDLOSS)," Proceedings of the Conference on Emergency Planning, Vol. 15, No. 1, La Jolla, CA: The Society for Computer Simulation, Jan. 1985.

This paper describes a computerized methodology to estimate earthquake damages, casualties, and dollar losses in the event of damaging earthquakes. The type of damage/loss estimates the FEDLOSS model provides are: (1) the expected physical damage caused by ground shaking, (2) the expected percentage of loss of function or usability, (3) the expected percentage of population killed and injured, and (4) expected losses due from collateral earthquake hazards such as ground failure, fire and inundation. The output of the FEDLOSS model can be used to drive other models such as the FEMA Earthquake Impacts Modeling System (FEIMS).

Nelson, W. R., "Reactor: An Expert System for Diagnosis and Treatment of Nuclear Reactor Accidents," Proceedings of the National Conference on Artificial Intelligence, AAAI-82, American Association for Artificial Intelligence, 1982.

REACTOR is an expert system under development at EG&G Idaho, Inc. that will assist operators in the diagnosis and treatment of nuclear reactor accidents. This paper covers the background of the nuclear industry and why expert system technology may prove valuable in the reactor control room. Some of the basic features of the REACTOR system are discussed, and future plans for validation and evaluation of REACTOR are presented. The concept of using both event-oriented and function-oriented strategies for accident diagnosis is discussed. The response tree concept for representing expert knowledge is also introduced.

Nelson, W. R., "Response Trees and Expert Systems for Nuclear Reactor Operations," Report No.: EGG-2293, Department of Energy, Washington, DC, Mar. 1984.

Ogawa, H., K. S. Fu, and J. T. P. Yao, "An Expert System for Structure Damage Assessment," Pattern Recognition Lett., Vol. 2, No. 6, Dec. 1984.

An expert system called SPERIL-II is introduced for the damage assessment of existing structures using the knowledge of experienced structural engineers. Fuzzy sets and Dempster and Shafer's Theory are used in this inexact inference method.

Petrie, J. F., and W. T. Hathaway, "Development of an Automated Emergency Response System (AERS) for Rail Transit Systems," Report No.: UMTA-MA-06-0152-84-4, U. S. Department of Transportation, UMTA, Office of Technical Assistance, Washington, D.C., Oct. 1984.

This report presents the results of an effort to deploy and evaluate automated emergency response systems (AERS). Developed initially by a train controller at the Bay Area Rapid Transit District (BART), the AERS is a computerized data bank containing equipment and facilities location information and predetermined response actions. Its purpose is to provide controllers, dispatchers and supervisors with a quick and accurate information retrieval system. In the development of UMTA's Recommended Emergency Preparedness Guidelines for rail transit systems, the AERS was identified as a decision-making aid that would be of value to the transit industry.

Russo, E. J., and R. R. Wilson, "Earthquake Vulnerability Analysis for Economic Impact Assessment," Proceedings of the Conference on Computer Simulation in Emergency Planning, La Jolla, CA: The Society for Computer Simulation, Jan. 1983.

This paper describes research to develop multidisciplinary techniques for a comprehensive damage and economic and industrial impact evaluation for catastrophic California earthquakes. The earthquake damage evaluation technique is patterned after damage evaluation techniques for nuclear war that have evolved over the past several decades. The economic impact evaluation techniques are based on a small area input-output modeling system in which the damage evaluations are used to generate evaluations of the economic dislocations. The damage evaluation technique integrates three components including geocoded seismic intensity simulations; geocoded inventories of housing, buildings, and facilities; and damage functions relating seismic intensity and structural characteristics of buildings and facilities to damage estimates. The damage evaluation technique is designed to simulate damages from hypothetical earthquakes and to provide rapid first damage estimates in the event of actual catastrophic earthquakes in California. The economic impact evaluation technique is designed to employ results of the damage evaluation methodology in conjunction with a recently developed interindustry capability. The new interindustry system provides for the evaluation of disaster economic impacts with an unprecedented geographic and economic sector detail. With this system, county, state, regional, and national economic impacts of catastrophic events can be evaluated for up to 470 economic sectors.

Wallace, W. A., and F. de Balogh, "Decision Support Systems for Disaster Management," Public Administration Review, Vol. 45, Special Issue, Jan. 1985.

This paper provides a conceptual framework for the employment of decision support systems (DSS) in the field of disaster management. This framework incorporates the following dimensions: (1) definition of DSS and its differences from MIS in the public sector, (2) stages in the life cycle of a disaster and the types of events occurring within stages, (3) a framework for conceptualizing decision making in disaster management at various levels -- operational, tactical, and strategic, (4) a matrix of information and modeling needs of planners and decision makers within each stage, and (5) selected examples of DSS applications and prototypes, developed on microcomputers, that serve to illustrate the role of DSS in the disaster management.

Waltz, D. L., "New Tools for Emergency Management," Proceedings of the First Symposium on the Theory and Application of Expert Systems in Emergency Management Operations, Special Publication 717, Gaithersburg, MD: National Bureau of Standards, 1986.

This paper explores a number of possible applications of AI to the problems of emergency management. First, it presents some far-future scenarios, illustrating the kinds of aid that AI could ultimately supply. While many of the applications will require many years of research before they can be realized, there are some near-term applications, especially for new natural language processing technologies; a few interesting possibilities are discussed. There are marked similarities between the problems of emergency management and the three key problems being addressed by the U.S. Department of Defense. These similarities are identified, and projections made about trickle-down effects of this DoD research. Finally, the paper discusses some obstacles to achieving the goals of automated emergency management systems and aids, and presents some ideas on how we can achieve these goals.

Williams, T. L., P. J. Organ, and C. L. Smith, "Diagnosis of Multiple Faults in a Nationwide Communications Network," Proceedings of the Eighth IJCAI, American Association for Artificial Intelligence, 1983.

The Network Diagnostic System (NDS) is an ARBY based expert system for fault isolation in a nationwide communications network (COMNET). Due to both the structure and function of the network, failures in COMNET are often multiple component failures (either dependent or independent) or intermittent failures. The maintenance procedure for isolating and correcting faults in COMNET exploits multiple types of knowledge, including the topological structure of COMNET, geographic organization, and frequency of failure information. Fault isolation in NDS is represented as a heuristic search through a space of hypotheses. The available diagnostic tests impose a refinement hierarchy on the space of hypotheses, enabling the exploitation of hierarchical search. Back links to more general hypotheses at higher levels in the refinement hierarchy are introduced to ensure the isolation and repair of multiple and intermittent failures. The NDS currently performs at the level of an intermediate COMNET diagnostician.

## APPENDIX A

## LIST OF JOURNAL AND PROCEEDINGS SOURCES


AAMSI Congress 83 Proceedings, Bethesda, MD: American Association of Medical Systems and Informatics, May 1983.

ACM SIGEUE Bulletin, Vol. 12, 1978.

Advances in Computers, Vol. 22, New York: Academic Press, 1983.

AI Magazine, Vol. 2, 1981; Vol. 3, 1982.

American Journal of Hospital Pharmacy, Vol. 37, 1980.

American Journal of Medicine, Vol. 60, 1976; Vol. 64, 1978.

American Scientist, Vol. 72, Sept.-Oct., 1984; Vol. 73, May-June 1985.

The American Statistician, Vol. 39, No. 1, Feb. 1985.

Annals of Mathematical Statistics, Vol. 37, 1967; Vol. 38, 1967.

Artificial Intelligence, Vol. 5, 1971; Vol. 8, 1974; Vol. 11, 1978; Vol. 12, 1979; Vol. 18, 1982; Vol. 20, 1983; Vol. 21, 1983; Vol. 22, 1984.

Australian Computer Journal, Vol. 12, No. 2, May 1980.

Automatica, Vol. 13, 1977.

Automedica, Vol. 4, No. 2/3, 1982.

Byte, Vol. 6, No. 9, Sept., 1981; Vol. 9, No. 5, May 1984, Vol. 10, No. 2, Jan. 1985; Vol. 10, No. 4, Apr. 1985.

Communications ACM 8, 1971.

Computer, Vol. 16, No. 2, Feb. 1983; Vol. 16, No. 10, Oct. 1983; Vol. 17, No. 3, March. 1984; Vol. 17, No. 9, Sept. 1984; Vol. 17, No. 10, Oct. 1984; Vol. 18, No. 1, Jan. 1985; Vol. 18, No. 4, Apr. 1985.

Computers and Mathematics, Vol. 9, No. 1, 1983; Vol. 11, No. 5, 1985.

The Computer Journal, Vol. 23, No. 4, 1980.

Computer Magazine, Feb. 21, 1985.

Computer Weekly, No. 876, Sept. 1983.

Computing Surveys, Vol. 12, 1980.

Computers and Biomedical Research, Vol. 15, No. 2, 1982.

Computerworld, Vol. 17, No. 20, May 1983; Vol. 17, No. 43, Oct. 1983; Vol. 18, No. 49A, Dec. 1984.

Creative Computing, Vol. 6, No. 3, Mar. 1980.

CRS Review, Nov./Dec. 1984.

CSELT Rapp. Tec., Vol. 11, No. 3, June 1983.

Cybernetics and Systems, Vol. 14, No. 2-4, Apr.-Dec. 1983.

Data Processing, Vol. 24, No. 3, Mar. 1982.

Datamation, Nov. 1, 1984.

Decision Sciences, Vol. 12, 1981.

Electronic Design, Vol. 32, No. 16, Aug. 1984; Vol. 33, No. 1, Jan. 1985.

EPRI Journal, June 1985.

European Journal of Operational Research, Vol. 10, 1982.

Expert Systems, Vol. 1, No. 2, Oct. 1984; Vol. 2, No. 1, Jan. 1985.

Fuzzy Sets and Systems, Vol. 11, 1983.

High Technology, Jan. 1984; Mar. 1985; Apr. 1985.

Human Systems Management, Vol. 4, No. 4, Autumn 1984.

ICL Technical Journal, 1980.

IEE Colloquium on Decision Support Aspects of Expert Sytems (Digest No. 67), London, England: IEE, May 1984.

IEEE Transactions on Circuits and Systems, Vol. CAS-22, No. 11, 1975.

IEEE Transactions on Man-Machine Systems, MMS-11, 1970.

IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 2, No. 5, Sept. 1980; Vol. 6, No. 1, Jan. 1984; Vol. 6, No. 2, Mar. 1984; Vol. 7, No. 3, May 1985.

IEEE Transactions on Systems, Management, and Cybernetics, Vol. SMC-12, No. 2, Mar./Apr. 1982; Vol. SMC-14, No. 6, Nov./Dec. 1984.

Information Age, Vol. 6, No. 4, Oct. 1984.

Information and Control, Vol. 8, 1965.

Information Sciences, Vol. 5, 1973.

_Interfaces_, Vol. 14, No. 2, Mar./Apr., 1984; Vol. 14, No. 3, May/June 1984; Vol. 14, No. 5, Sept./Oct. 1984.

_International Journal of Man-Machine Studies_, Vol. 6, 1974; Vol. 8, 1976; Vol. 9, 1977; Vol. 10, 1978; Vol. 12, 1980; Vol. 18, 1983; Vol. 19, 1983; Vol. 20, 1984; Vol. 21, 1984.

_International Journal of Policy Analysis and Information Systems_, Vol. 4, No. 4, Dec. 1980.

_Journal of American Planning Association_, Vol. 3, July 1981.

_Journal of American Society for Information Science_, Vol. 35, No. 5, Sept. 1984.

_Journal of Hazardous Material_, Special Issue, Vol. 4, 1981.

_Journal of Information Systems_, Vol. 1, No. 2, Fall 1984.

_Journal of Mathematical Biology_, Vol. 1, No. 1, 1974.

_Journal of Operational Research Society_, Vol. 36, No. 2, Feb. 1985.

_Journal of the Royal Statistical Society_, Series B., Vol. 25, 1963; Series B, Vol. 30, No. 2, 1968.

_Journal Research Incorporated Marconi Review_, Vol. 1, No. 1, 1983.

_Machine Intelligence 6_, Chichester, England: Horwood, 1971.

_Machine Intelligence 9_, Chichester, England: Horwood, 1981.

_Machine Intelligence 10_, Chichester, England: Horwood, 1982.

_Management Science_, Vol. 17, No. 4, Dec. 1970; Vol. 29, No. 1, Jan. 1983.

_Mathematical Biosciences_, Vol. 23, 1975.

_Medical Information_, Vol. 8, No. 3, July/Sept., 1983; Vol. 9, No. 1, Jan./Mar. 1984.

_Methods of Information in Medicine_, Vol. 21, No. 3, July 1982.

_Microsystems_, Aug. 1983, Jan. 1984.

_National Online Meeting Proceedings_, Apr. 1984.

_New Generation Computers_, Vol. 1, No. 1, 1983; Vol. 2, No. 2, 1984.

_Pattern Recognition Letters_, Vol. 2, No. 6, Dec. 1984.

_Personal Computer World_, Vol. 7, No. 6, June 1984.

_Practical Computer_, Vol. 7, No. 10, Oct. 1984.

Proceedings of the AAAS Selected Symposium 54 on Mathematical Frontiers of the Social and Policy Sciences, Washington, DC; American Association for the Advancement of Science, 1981.

Proceedings of the AFIPS National Computer Conference, Arlington, VA: AFIPS Press, Vol. 45, 1976; Vol. 51, 1982.

Proceedings of the 1984 American Control Conference, Vol. 1, New York: IEEE, June 1984.

Proceedings of the ASME Symposium on Engineering Databases: Software for On-Line Applications, San Antonio, TX: ASME, June 1984.

Proceedings of the Autotestcon '82: the 1982 IEEE International Automatic Testing Conference, New York: IEEE, Oct. 1982.

Proceedings of the 1984 Carnahan Conference on Security Technology, New York: IEEE, May 1984.

Proceedings of the Colloquium on Application of Knowledge Based (or Expert) Systems, London, England, 1982.

Proceedings of the COMPAS '84, Berlin, Germany: VDE-Verlag, Oct. 1984.

Proceedings of the COMPCON 83 Fall, Silver Spring, MD: Computer Society Press, Oct. 1983.

Proceedings of the COMPSAC 80, Oct. 1980.

Proceedings of the Conference on Computer Simulation in Emergency Planning, La Jolla, CA: The Society for Computer Simulation, Vol. 11. No. 2, Jan. 1983.

Proceedings of the Conference on Emergency Planning, La Jolla, CA: The Society for Computer Simulation, Vol. 15, No. 1, Jan. 1985.

Proceedings of the Conference on Software Development Tools, Techniques, and Alternatives, Silver Spring, MD: IEEE Computer Society Press, July 1983.

Proceedings of the ECAI-82, 1982.

Proceedings of the Eighth IJCAI, American Association for Artificial Intelligence, 1983.

Proceedings of the Eleventh Collegium Internationale Neuro-Psychopharmacologicum, Vienna, 1978.

Proceedings of the Ergonomics Society's Conference, London, England: Taylor and Francis, Mar. 1983.

Proceedings of the Fifteeenth Hawaii International Conference on System Sciences, Honolulu, Hawaii: Hawaii International Conference of System Sciences, Vol. 2, 1983.

Proceedings of the Fifteenth Pittsburgh Modeling and Simulation Conference, Pittsburgh, PA, Aug. 1984.

Proceedings of the First Annual Symposium on Computer Application in Medical Care, New York: IEEE, Oct. 1977.

Proceedings of the First Symposium on the Theory and Application of Expert Systems in Emergency Management Operations, Special Publication 717, Gaithersburg, MD: National Bureau of Standards, 1986.

Proceedings of the Fourth MIT/ONR Workshop on Distributed Information and Decision Systems, 1982.

Proceedings of the Fourteenth International Conference on Fault-Tolerant Computing, Silver Spring, MD: IEEE Computer Society Press, 1984.

Proceedings of the Fourteenth Symposium on the Interface of Computer Science and Statistics, New York: Springer-Verlag, 1982.

Proceedings of the Fourth Annual Symposium on Computer Applications in Medical Care, 1980.

Proceedings of the Fourth Congress on Pattern Recognition and Artificial Intelligence, Le Chesnay, France: Inst. Nat. Recherche Inf. and Autom., 1984.

Proceedings of the IEEE, Vol. 67, No. 9, Sept. 1979; Vol. 71, No. 11, Nov. 1983; Vol. 72, No. 12, Dec. 1984.

Proceedings of the IEEE Computer Society Conference on Pattern Recognition and Image Processing, New York: IEEE, 1981.

Proceedings of the 1983 IEEE Computer Society Workshop on Computer Architecture for Pattern Analysis and Image Database Management, Silver Spring, MD: IEEE Computer Society Press, Oct. 1983.

Proceedings of the IEEE International Conference on Computer Design: VLSI in Computers, New York: IEEE, 1983.

Proceedings of the 1983 IEEE Military Communications Conference, New York: IEEE, Vol. 3, 1983.

Proceedings of the 1984 IEEE National Aerospace and Electronics Conference (NAECON), New York: IEEE, Vol. 2, May 1984.

Proceedings of the IFIP-IMIA Working Conference, Amsterdam, Netherlands: North-Holland, 1983.

Proceedings of the International Conference on Cybernetics and Society, New York: IEEE, Oct. 1978; Oct. 1980; Oct. 1981; Oct. 1982.

Proceedings of the International Conference on Data Engineering, Silver Spring, MD: IEEE Computer Society Press, Apr. 1984.

Proceedings of the International Conference on Data Engineering, Silver Spring, MD: IEEE Computer Society Press, Apr. 1984.

Proceedings of the Joint BCS and ACM Symposium on Research and Development in Information Retrieval, Cambridge, England: Cambridge Univ.: 1984.

Proceedings of the Joint Services Workshop on Artificial Intelligence in Maintenance, Boulder, CO, Oct. 1983.

Proceedings of the Micro-Delcon '81, The Delaware Bay Computer Conference, Mar. 1982.

Proceedings of the National Conference on Artificial Intelligence, American Association for Artificial Intelligence, AAAI-80, 1980; AAAI-82, 1982.

Proceedings of the Seventh IJCAI, American Association for Artificial Intelligence, 1981.

Proceedings of the Seventh International Conference on Software Engineering, New York: IEEE, Mar. 1984.

Proceedings of the Sixteenth Annual Meeting of the American Institute for Decision Sciences, Atlanta, GA: American Inst. for Decision Sciences, Nov. 1984.

Proceedings of the Sixth Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, Baltimore, MD: ACM, June 1983.

Proceedings of the Sixth IJCAI, American Association for Artificial Intelligence, 1979.

Proceedings of the SPIE International Society of Optical Engineering, Vol. 485, May 1984.

Proceedings of the Statistical Computing Section, American Statistical Association, 1983.

Proceedings of the Third Annual International Phoenix Conference on Computers and Communications, Silver Spring, MD: IEEE Computer Society Press, 1984.

Proceedings of the Third IJCAI, American Association of Artificial Intelligence, 1973.

Proceedings of the Third International Conference on Information Systems, Ann Arbor, MI, 1982.

Proceedings of the Third National Conference and Exhibition on Computers in Personnel: Making Manpower Profitable, London, England: Institute of Personnel Management, June 1984.

Proceedings of the Thirteenth International Symposium on Mutiple-Valued Logic, New York: IEEE, May 1983.

_Proceedings of Trends and Applications 1983 on Automating Intelligent Behavior: Applications and Frontiers_, Silver Spring, MD: IEEE Computer Society Press, May 1983.

_Proceedings of the Twenty-first Annual Meeting of the Association for Computational Linguistics_, Menlo Park, CA: Association of Computational Linguistics, June 1983.

_Proceedings of Videotex '84 International_, London, England: Online Publications, 1984.

_Proceedings of the Workshop on Computer Vision: Representation and Control_, Silver Spring, MD: IEEE Computer Society Press, 1984.

_Public Administration Review_, Vol. 45, Special Issue, Jan. 1985.

_Science_, Vol. 197, Sept. 1977; Vol 217, Sept. 1982; Vol. 220, Apr. 1983; Vol. 223, Mar. 1984; Vol. 228, May 1985; Vol. 228, Apr. 1985.

_Scientific American_, Vol. 247, No. 4, Oct. 1982.

_Society of Science and Information Studies_, Vol. 5, No. 1, Jan. 1985.

_SIGART Newsletter_, No. 87, Jan. 1984.

_Systems International_, Vol. 12, No. 8, Aug. 1984.

_Technology Review_, Nov./Dec., 1983.

_TIMS Studies in the Management Sciences_, Vol. 20, 1984.

_Transactions in Programming and Computer Software_, Vol. 9, No. 4, July-Aug. 1983.

_Water Resources Research_, Vol. 19, No. 6, Dec. 1983.

_VAX/RSTS Prof._, Vol. 6, No. 5, Oct. 1984.

## APPENDIX B

## GLOSSARY OF TECHNICAL TERMS

Algorithm - A well-defined computational procedure for solving a problem in a finite number of steps.

AND/OR Graph - A tree-like structure with two types of nodes: AND nodes for which several successors of a node have to be accomplished (or considered), and OR nodes for which only one of several of the node's successors are to be considered.

Artificial Intelligence - A discipline devoted to developing and applying computational approaches to intelligent behavior.

Atom - An indivisible element. A logical proposition that cannot be broken down into other propositions.

Backtracking - Describes a return to an earlier point in a search space, usually due to a depth-first search failure. Also a name given to depth-first backward reasoning.

Backward Chaining - Describes a system that begins a solution search by examining a limited set of production rules with defined goals as its starting point.

Bidirectional Search - A solution search procedure which proceeds simultaneously from both the start node and from a set of goal nodes.

Blackboard - A common working data storage with which the various system elements communicate with each other.

Breadth-First - A search technique in which, starting with the root node, the nodes in the search tree are generated and examined level by level before proceeding deeper. This approach is guaranteed to find an optimal solution if it exists.

Combinatorial Explosion - Describes the rapid growth of possibilities as the search space expands. If each branch point (decision point) has an average of n branches, the search space tends to expand as $n^d$, as the depth of search, d, increases.

Computer Architecture - Describes the way in which the various computational elements are interconnected to produce a computational function.

Computer Vision - The symbolic description of a scene depicted in an image developed by a computer using visual sensory inputs such as photographs.

Connectives - Describes the operators (such as AND, OR, etc.) connecting statements in logic so that the truth-value of the composite is determined by the truth-value of the components.

Consistency Enforcer - An element in expert systems responsible for maintaining consistent representation of the emerging solution.

Controller - Refers to a distinct entity in an expert system, consisting of an interpreter, a scheduler and a consistency enforcer.

Control Structure - That aspect of an expert system which determines how to use the rules contained in the knowledge base.

Data Base - Describe an organized collection of information on some subject.

Data-Driven - A problem solving technique which uses forward reasoning.

Declarative Knowledge - Facts and assertions about a problem that is represented in the global data base.

Depth-First - A search technique in which one proceeds from the root node to one of the successor nodes and then to one of that node's successor nodes, etc., until a solution is reached or the search is forced to backtrack.

Domain - The range of the problem area of interest.

Event-Driven - A problem solving procedure that uses forward-chaining driven by the current status of the problem.

First-Order Predicate Calculus - A system of logic in which variables representing predicates are not admitted.

Forward Chaining - Describes a system that starts with a subset of evidence and proceeds to invoke the production rules in a forward direction, continuing until no further production rules can be invoked.

Frame - A knowledge structure (or database) used to describe the attributes that an object possesses.

Generate and Test - A problem solver consisting of two basic modules: one that enumerates (generates) possible solutions and the other that evaluates (tests) each proposed solution, either accepting it or rejecting it.

Global Data Base - The name given to the data base containing information on the specific problem, its status, and the solution process, and is accessible to all the components of an expert system.

Goal Driven - A problem solving technique that uses the goal of the problem to guide the solution search using backward chaining.

Heuristic - A technique that human experts use to improve the efficiency of a problem-solving process, even though its use cannot be justifiied theoretically.

Inference Engine -That part of an expert system responsible for applying the appropriate rules and stopping when a termination condition has been met; also called the "control structure."

Intelligent Front-End - Used to describe that aspect of the user interface which makes a system more flexible and user-friendly.

Interface - The system by which a user interacts with the computer; in general, the connection between two components.

Interpreter - That element of the controller which executes rules according to an agenda.

Invoke - To put a rule into action, usually by satisfying a precondition.

Irrevocable Control Strategy - A control strategy in which an applicable rule once applied is not allowed to be reconsidered later.

Justification - A statement about the line of reasoning used by an expert system in drawing a conclusion or for demanding some additional piece(s) of information.

Justifier - An element of the expert system that performs the justification task.

Knowledge Base - The collection of facts and rules resident in the data base of an expert system.

Knowledge Engineer - A person whose task is to extract knowledge from an expert and represent that knowledge in a computer.

Logical Representation - Knowledge representation using logical formulas that capture a partial description of the world.

Meta-Knowledge - The knowledge which allows a system to reason about actions, processes and plans.

Meta-Problems - These are higher level problems involving a system's own actions, processes and plans.

Meta-Rules - Higher level rules used to reason about lower level rules.

Minimaxing - A search procedure in which a system determines the score at each node in a game tree by either picking the minimum or maximum scores.

Model Driven - A problem solving technique that uses top-down approach and generates inferences to be verified based on the domain model present in the system.

Modus Ponens - A logical rule of inference which has the form:

If $E_1$ is true, then $E_2$ is true.
$E_1$ is true.
Therefore $E_2$ is true.

Modus Tolens - A logical rule of inference which has the form:
If $E_1$ implies $E_2$, then not $E_1$ implies not $E_2$.

Natural Language - A term used to refer to languages like English.

Parallel Processing - The simultaneous processing of data, as opposed to the usual sequential processing done by conventional computers.

Parser - Refers to the syntactic program that segments sentences into meaningful phrases.

Pattern Recognition - A term given to describe the task of classifying a set of data into predetermined categories.

Predicate Calculus - A modification of propositional calculus to allow the use of variables and functions of variables.

Problem State - The condition of a problem at any given instant.

Procedural Knowledge - The knowledge about a problem that is represented in the rules.

Production Rules - Rules used in production systems, that have the general form:

If: Logical conditions are satisfied
Then: Take the indicated action.

Production System - A system based on a collection of production rules.

Propositional Calculus - An elementary logic that uses argument forms to deduce the truth or falsehood of a new proposition from known propositions.

Rule-Interpreter - see Interpreter.

Rules - Elements in the knowledge base of an expert system that tell the system which direction to take or what to do next.

Scheduler - That element in the control structure that maintains control of the agenda and determines which pending action should be executed next.

Script - A frame-like knowledge structure used to represent related sequence of events.

Search Space - The implicit graph representing all the possible states of the system that may have to be searched to find a solution. In many cases the search space is infinite. The term search space is also used for non-state-space representations.

Semantic Net - A directed graph of nodes and connecting links. The nodes represent objects, events, concepts, situations or actions and the links represent properties and relationships between them.

Solution Path - A successful path through a search space.

Speech Recognition - Recognition by a computer of spoken words or sentences.

State Space - A space in which the nodes represent the system state and the connecting arcs represent operators that can be used to transform the state from which the arcs emanate to the state at which they arrive.

Symbolic Reasoning - Use of abstract representations or symbols (for concrete objects) and reasoning methods in arriving at the goal state from the initial state, or vice versa.

Top-Down Approach - A problem solving technique which is goal directed or expectation-guided based on models or other knowledge; also referred to as "hypothesize and test" approach.

Truth Value - One of the two values associated with propositions in logic - True or False.

User Interface - see "interface."

APPENDIX C

SOME COMMON ABBREVIATIONS

AAAI - American Association for Artificial Intelligence

ACE (expert system) - Automated Cable Expertise

AERS - Automated Emergency Response System

AGE (expert system) - Attempt to Generalize

AI - Artificial Intelligence

ART (shell) - Automatic Reasoning Tool

CAI - Computer Assisted Instruction

CASNET (expert system) - Causal Associated Network

CF - Certainty Factor

COMEX (expert system) - Compact Knowledge-Based Expert Systems

COMNET - Communications Network

CONGEN (expert system) - Constrained Generator

CRIB (expert system) - Computer Retrieval Incidence Bank

DSS - Decision Support Systems

ES - Expert System

IEMIS - Integrated Emergency Management and Information System

IFE - Intelligent Front End

IJCAI - International Joint Conference on Artificial Intelligence

KAS (expert system) - Keyword Access System

KB - Knowledge Base

KEE (shell) - Knowledge Engineering Environment

KNOBS (expert system) - Knowledge Based System

LISP - List Processor

MB - Measure of Belief

MCDM - Multiple Criteria Decision Making

MD - Measure of Disbelief

MLSE - Multi-Layered Software Environment

NAECON - National Aerospace and Electronics Conference

NLI - Natural Language Interface

PR - Production Rule

PROLOG - Programming in Logic

RLL (shell) - Representation Language Language

RITA (shell) - Rule-directed Interactive Transaction Agent

ROSIE (shell) - Rule Oriented System for Implementing Expertise

SACON (expert system) - Structural analysis consultant

TATR (expert system) - Tactical air target recommender

TIMM (shell) - The Intelligent Machine Model

VLSI - Very Large Scale Integrated Circuits

NBS-114A (REV. 2-80)

| U.S. DEPT. OF COMM. BIBLIOGRAPHIC DATA SHEET (See instructions) | 1. PUBLICATION OR REPORT NO. NBS/SP-728 | 2. Performing Organ. Report No. | 3. Publication Date November 1986 |
|---|---|---|---|

4. TITLE AND SUBTITLE

Expert Systems and Emergency Management: An Annotated Bibliography

5. AUTHOR(S) Saul I. Gass, Suneel Bhasker, and Robert E. Chapman

6. PERFORMING ORGANIZATION (If joint or other than NBS, see instructions)

NATIONAL BUREAU OF STANDARDS
DEPARTMENT OF COMMERCE
Gaithersburg, MD 20899

7. Contract/Grant No.

8. Type of Report & Period Covered

Final

9. SPONSORING ORGANIZATION NAME AND COMPLETE ADDRESS (Street, City, State, ZIP)

Federal Emergency Management Agency
500 C Street, SW
Washington, DC 20472

10. SUPPLEMENTARY NOTES

Library of Congress Catalog Card Number 86-600591

☐ Document describes a computer program; SF-185, FIPS Software Summary, is attached.

11. ABSTRACT (A 200-word or less factual summary of most significant information. If document includes a significant bibliography or literature survey, mention it here)

This report is the result of an in-depth review of the recent technical literature on expert systems. The material contained in this report provided a basis for assessing the potential for using expert systems in emergency management operations. In choosing the material for inclusion in this report, special emphasis was placed on those aspects of expert systems which addressed the types of problems encountered in emergency management operations.

The report is designed for use as a resource document and as a tutorial on expert systems and emergency management. Each chapter consists of a brief topic essay followed by a set of references which expand on the main themes of the essay.

12. KEY WORDS (Six to twelve entries; alphabetical order; capitalize only proper names; and separate key words by semicolons)

artificial intelligence; emergency management; expert systems

13. AVAILABILITY

[XX] Unlimited
☐ For Official Distribution. Do Not Release to NTIS
[XX] Order From Superintendent of Documents, U.S. Government Printing Office, Washington, D.C. 20402.
☐ Order From National Technical Information Service (NTIS), Springfield, VA. 22161

14. NO. OF PRINTED PAGES
179

15. Price

USCOMM-DC 6043-P80

*U.S. Government Printing Office: 1986 0-181-076 (40194)

# NBS Technical Publications

## Periodical

**Journal of Research**—The Journal of Research of the National Bureau of Standards reports NBS research and development in those disciplines of the physical and engineering sciences in which the Bureau is active. These include physics, chemistry, engineering, mathematics, and computer sciences. Papers cover a broad range of subjects, with major emphasis on measurement methodology and the basic technology underlying standardization. Also included from time to time are survey articles on topics closely related to the Bureau's technical and scientific programs. Issued six times a year.

## Nonperiodicals

**Monographs**—Major contributions to the technical literature on various subjects related to the Bureau's scientific and technical activities.

**Handbooks**—Recommended codes of engineering and industrial practice (including safety codes) developed in cooperation with interested industries, professional organizations, and regulatory bodies.

**Special Publications**—Include proceedings of conferences sponsored by NBS, NBS annual reports, and other special publications appropriate to this grouping such as wall charts, pocket cards, and bibliographies.

**Applied Mathematics Series**—Mathematical tables, manuals, and studies of special interest to physicists, engineers, chemists, biologists, mathematicians, computer programmers, and others engaged in scientific and technical work.

**National Standard Reference Data Series**—Provides quantitative data on the physical and chemical properties of materials, compiled from the world's literature and critically evaluated. Developed under a worldwide program coordinated by NBS under the authority of the National Standard Data Act (Public Law 90-396). NOTE: The Journal of Physical and Chemical Reference Data (JPCRD) is published quarterly for NBS by the American Chemical Society (ACS) and the American Institute of Physics (AIP). Subscriptions, reprints, and supplements are available from ACS, 1155 Sixteenth St., NW, Washington, DC 20056.

**Building Science Series**—Disseminates technical information developed at the Bureau on building materials, components, systems, and whole structures. The series presents research results, test methods, and performance criteria related to the structural and environmental functions and the durability and safety characteristics of building elements and systems.

**Technical Notes**—Studies or reports which are complete in themselves but restrictive in their treatment of a subject. Analogous to monographs but not so comprehensive in scope or definitive in treatment of the subject area. Often serve as a vehicle for final reports of work performed at NBS under the sponsorship of other government agencies.

**Voluntary Product Standards**—Developed under procedures published by the Department of Commerce in Part 10, Title 15, of the Code of Federal Regulations. The standards establish nationally recognized requirements for products, and provide all concerned interests with a basis for common understanding of the characteristics of the products. NBS administers this program as a supplement to the activities of the private sector standardizing organizations.

**Consumer Information Series**—Practical information, based on NBS research and experience, covering areas of interest to the consumer. Easily understandable language and illustrations provide useful background knowledge for shopping in today's technological marketplace.

*Order the* **above** *NBS publications from: Superintendent of Documents, Government Printing Office, Washington, DC 20402.*

*Order the* **following** *NBS publications—FIPS and NBSIR's—from the National Technical Information Service, Springfield, VA 22161.*

**Federal Information Processing Standards Publications (FIPS PUB)**—Publications in this series collectively constitute the Federal Information Processing Standards Register. The Register serves as the official source of information in the Federal Government regarding standards issued by NBS pursuant to the Federal Property and Administrative Services Act of 1949 as amended, Public Law 89-306 (79 Stat. 1127), and as implemented by Executive Order 11717 (38 FR 12315, dated May 11, 1973) and Part 6 of Title 15 CFR (Code of Federal Regulations).

**NBS Interagency Reports (NBSIR)**—A special series of interim or final reports on work performed by NBS for outside sponsors (both government and non-government). In general, initial distribution is handled by the sponsor; public distribution is by the National Technical Information Service, Springfield, VA 22161, in paper copy or microfiche form.

**U.S. Department of Commerce**
National Bureau of Standards
Gaithersburg, MD 20899