Artificial Intelligence and Expert Systems

INPUT®

ARTIFICIAL INTELLIGENCE AND EXPERT SYSTEMS

ARTIFICIAL INTELLIGENCE AND EXPERT SYSTEMS

CONTENTS

EXHIBITS

I INTRODUCTION

- The fascination with artificial intelligence (AI) predates the electronic computer. In the 1800s a mechanical chess-playing robot was developed and enjoyed some brief success until it was discovered that a midget was concealed inside the mechanism. Although advances in electronic technology have permitted increasing sophistication in hiding the human in the mechanism, there is still a substantial carnival element of conning the rubes associated with the application of computer technology--especially in the adaptation of terminology. As one software manager was quoted in a recent issue of the Applied Artificial Intelligence Reporter: "Before long, anything that contains a slick algorithm will be declared to possess AI."

- INPUT's poll to determine what would be covered in this report was unanimous on two points. Users want definitions and an understanding of what constitutes AI--specifically, of how expert, or knowledge-based, systems are "different" from conventional data processing systems.

- In addition, several very pragmatic questions were raised such as: "How much is the trend toward expert systems going to cost me?" and "How can the cost of expert systems be justified?" Both are very good questions, because most of the current excitement is being generated not by users but by the media.

- The purpose of this study is to answer these questions with special emphasis upon anticipated problems. Areas to be covered include:

- Definitions of AI and expert systems.

- AI and operations research (OR) affinities.

- Hardware, software, and human considerations--including costs.

- Quality assurance (QA) implications.

- Explanation of applications and domains.

- Integration of information systems activities.

- Obviously, with such a wide range of topics it will be impossible to explore specific areas in any detail. In fact, to do so would be misleading, because substantial controversy remains concerning many major issues.

- You will find that this report has more of an academic orientation than most INPUT reports. This is true not only because AI is in the process of emerging from an academic environment, but also because there are many fundamental technical and philosophical questions about AI that remain unanswered.

- The methodology employed in developing this report was to review AI information, to talk with knowledgeable people, and to reach conclusions intuitively. In other words, this product contains all natural intelligence without any artificial flavoring, preservatives, or supplements. Taken with reasonable exercise of your own natural intelligence, it should benefit both your personal and your corporate health and welfare.

II EXECUTIVE SUMMARY

- This executive summary is designed in presentation format in order to:

  - Help the reader review the key research findings.

  - Provide a general presentation and script to facilitate group communications.

- The key points of the entire report are summarized in Exhibits II-1 through II-5. On the left-hand page facing each exhibit is a script explaining the exhibit's contents.

- The entire field of artificial intelligence is shrouded in controversy, and research and development is bounded by the loading edge of many disciplines, including computer science. In many ways this report is like the first version of an expert system, and the executive summary represents only the surface of a complex subject. The reader is expected to understand only the bounds of the AI domain.

A. ATTITUDES TOWARD NEW TECHNOLOGIES

- When Charles Babbage conceived the analytical engine, the Countess of Lovelace (who interpreted Babbage's work) observed that there was a natural tendency to first overrate and then undervalue all such developments. This is perhaps the most profound observation ever made concerning computer technology.

- Consider all the solutions and final answers to systems development. They include COBOL, Data Base Management Systems, Structured Programming, Fourth-Generation Languages, etc. All have gone, or are going, through some phase of Lady Lovelace's cycle.

- It is certain that artificial intelligence and expert systems will be subject to the same cycle, but uncertainty arises in determining exactly where we are in the cycle.

- Expert systems, unlike traditional data processing systems, deal with uncertainty--with those "maybe...I wonder" problems that have no specific solution--rather than the straightforward simplicity of accounting type problems, which lead themselves to yes or no solutions.

- Confronted with uncertainty, it is certain that both humans and expert systems will reach varying solutions to the same problem. Establishing probabilities associated with the quality of such solutions is a difficult problem.

EXHIBIT II-1

INPUT®

ATTITUDES TOWARD NEW TECHNOLOGIES

The Lovelace Cycle -- New Technology Is:

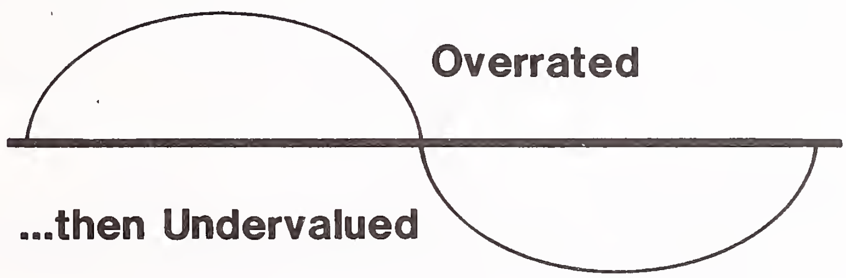

The Solution to Uncertainty Is:

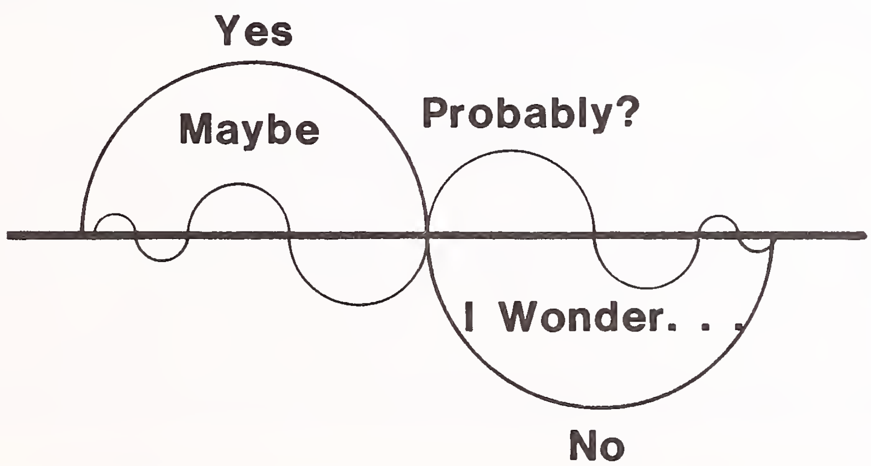

B. DATA, INFORMATION, AND KNOWLEDGE

- Data are facts and a great deal of progress has been made in storing data on electronic media for processing and for access by computers.

- Information is the result of analysis and may be based on facts and/or opinions. It is normally stored on paper--in file cabinets or libraries. Office automation is concerned primarily with the movement of information from paper to electronic media where it can be more readily processed, formatted, and accessed. Despite some progress, most information continues to be stored on paper.

- Knowledge is the result of human analysis of both data and information. The process of human knowledge acquisition (learning) and storage (memory) are not very well understood. Confronted with a vast array of data, information, and uncertainty, humans have the ability to solve problems (and make decisions) with varying degrees of success.

- At present, computer information systems facilitate making data and information more readily available to human beings, but human knowledge is required for practically all complex problem solving and decision making. Expert systems are attempting to put human knowledge in electronic form (knowledge bases), and the process of doing this is referred to as knowledge engineering. (This is represented by the broken-line circle on the exhibit.)

- Conceptually, data, information, and knowledge would converge into intelligent electronic systems, but progress has been excruciatingly slow. Despite considerable publicity, there are significant barriers to the development of intelligent machines for practical problem solving.

DATA, INFORMATION, AND KNOWLEDGE

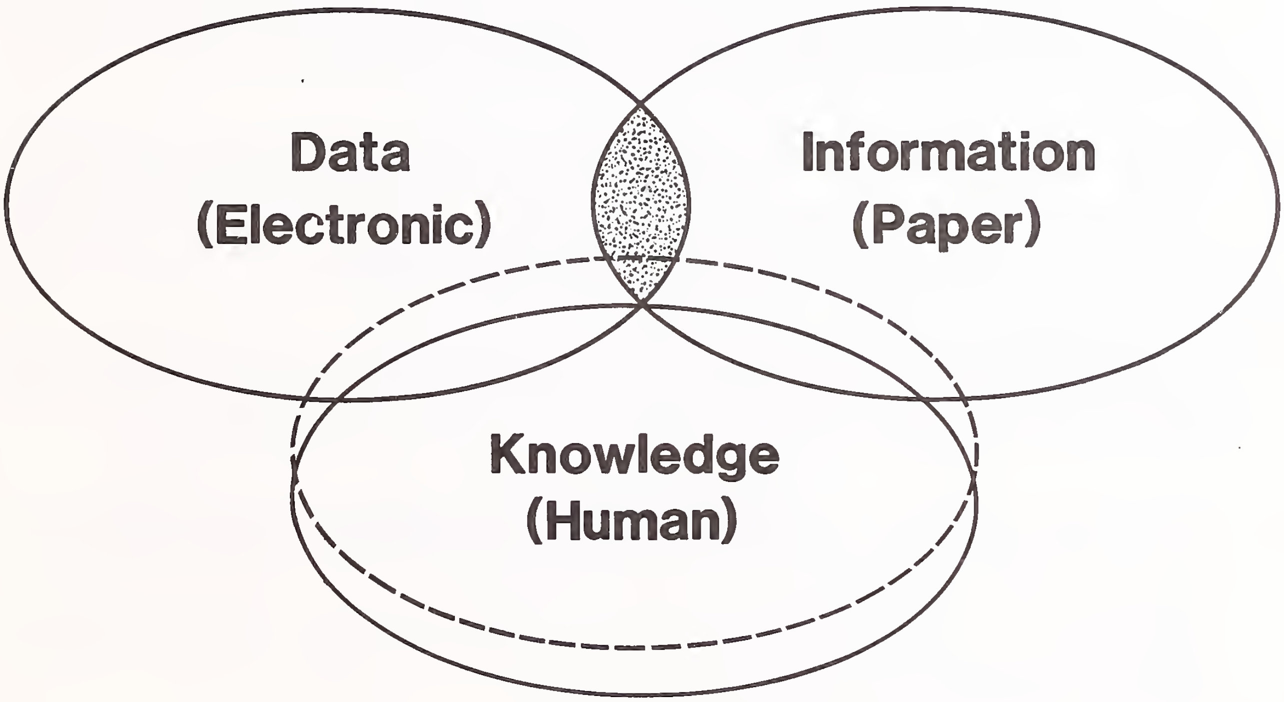

C. DON'T UNDERRATE PROBLEMS

- The primary reason the Lovelace cycle exists is because the computer industry has traditionally underrated implementation problems. As artificial intelligence and expert systems emerge from the academic environment, there are indications that implementation problems will be especially severe.

- There is fundamental ignorance concerning how humans learn, how knowledge is stored, and how intuition affects complex problem solving. These difficult problems are being disregarded by many AI practitioners.

- The AI community has isolated itself from many practical solutions developed in operations research and from much of the knowledge that has been gained from implementation of data base and information systems.

- A new profession of knowledge engineering is being created, without the benefit of the experience gained from complex systems analysis.

- Current expert systems lack flexibility and adaptability, and they are inclined to fail as they approach the periphery of these narrow domains (areas of expertise). They do not know the limits of their knowledge and cannot adjust to the knowledge level of the human expert.

- Current experience indicates that excessive maintenance is required to ensure acceptable performance and systems quality.

EXHIBIT II-3

INPUT®

DON'T UNDERRATE PROBLEMS

- Fundamental Ignorance About Human Knowledge Acquisition and Memory
- Poor Interdisciplinary Communications
- Inexperienced Personnel
- Narrow Domains (Applications) of Expert Systems
- Poor Quality of Systems

D. DON'T UNDERVALUE A.I. CONCEPTS

- As information systems become more complex, it becomes essential to formalize how data and information are being used--in other words, how they are expected to contribute to human knowledge.

  - End-user access to large central data bases makes them subject to misinterpretation and to the generation of conflicting information.

  - As more information becomes available, there is a tendency for quality (or the ability to assimilate) to diminish, and the acquisition of knowledge becomes more complex.

  - Individuals using available data and information are using personal computers and human knowledge to generate more information by developing their own systems. It becomes essential to understand the "human element" applied to this new information.

- It becomes essential to apply advanced tools from operations research, statistics, and AI if decision support systems are to be effective in such a complex environment. In other words, there is tremendous uncertainty in the emerging environment.

- Intelligent systems are required--especially in the sense of being able to explain how and why they reached the conclusion (or generated the information) they did.

- As systems become more complex, it becomes increasingly difficult to make them "user friendly." Intelligent systems are required to guide and educate users.

INPUT®

DON'T UNDERVALUE A.I. CONCEPTS

- I.S. Complexity Requires Formalization of Knowledge
  - Large Data Bases
  - Information Explosion
  - Distributed Systems Development
- New Tools Are Needed
- Intelligent Systems Are Required:
  - To Justify Conclusions
  - To Educate Users

E. A POSITIVE COURSE OF ACTION

- An advanced technology research function should be established under the IS department. AI research has had a narrow focus in its application to business problems. Knowledge-based development is where data base development was twenty years ago, and it is more complex. Basic knowledge concerning OR, statistics, information theory, decision theory, and library science will be required in addition to AI.

- There are few practical applications of expert systems at present, and the narrowness of their domains dictates that great care be exercised in exploring the potential of expert systems. This selection problem is compounded by the extremely broad array of "potential applications" currently being publicized. A non-critical problem would probably be best.

- Tools such as LISP machines are currently dictating the problems being addressed by AI. LISP is an old language that has never been formalized (there is practically an infinite variety of AI languages). General claims of systems development productivity should be viewed with a high degree of skepticism.

- Management should concentrate on quality and cost. Do not apply AI in order to be on the "leading edge," especially when better solutions are available. Expert systems are going to be costly in terms of both machine and human resources. In addition, they may not work, and you may not know it.

EXHIBIT II-5

INPUT®

A POSITIVE COURSE OF ACTION

- Establish a Research Function
  - OR
  - Statistics
  - AI
  - Information Theory
  - Decision Theory
  - Library Science
- Select Expert Systems Applications Carefully
- Evaluate AI Tools
- Concentrate on Cost and Quality

III WHAT IS ARTIFICIAL INTELLIGENCE?

A. A BRIEF HISTORY

- In interpreting the work of Charles Babbage on the analytical engine, Ada Augusta, Countess of Lovelace, stated that:

  - It is good to be wary of exaggerating the ideas that arise from the process of the analytical engine. There is often a tendency, in considering any new topic, to initially overrate the technology, by emphasizing the interesting or amazing aspects of it. Then, when we realize it doesn't meet our expectations, we tend to undervalue the true condition of the technology.

  - The analytical engine does not pretend to originate anything. The device can perform only what it is ordered to perform. It can follow analysis but cannot foresee any analytical relations or truths. The purpose and capacity of the device is to help make available what we are aware of already.

- Lady Lovelace's sage observations are as appropriate today as they were over 100 years ago. The tendency to overrate and then undervalue can clearly be seen whether we consider the "electronic brains" of the 1950s, the higher level languages in the 1960s, data base management systems in the 1970s, or decision support systems in the 1980s. AI and expert systems will be subject to the same tendency, and if Lady Lovelace's advice can help us moderate this tendency, she deserves more recognition than having a programming language named after her.

- However, the fascination with a "thinking machine" simply will not go away despite the clearly stated lack of "pretensions" on behalf of the analytical engine. Alan Turing (Computing Machinery and Artificial Intelligence) went to some length to present negative views in his exploration of whether machines could think. After devising the "Turing test" to resolve that question, and after concluding that a machine could be constructed to pass that test, he specifically addressed Lady Lovelace's analysis of the analytical engine's capability--specifically, of the implication that a machine cannot really do anything new. Turing responded directly to Lady Lovelace's "objection."

    - Turing provided some explanation for the viewpoint that machines cannot do anything new. He suggested that thinkers often erroneously assume that the mind assimilates facts together with the consequences that arise from such facts. (It is an assumption, he claims, that philosophers and mathematicians are subject to, but that is, nevertheless, fallacious.) The significance is that this assumption undervalues the natural development of consequences arising from facts and ideas.

    - It should be pointed out that Turing did not see fit to even discuss the question of "can machines think?" because he considered it meaningless. He conjectured that by the end of this century one would be able to speak of machines thinking without expecting to be contradicted. He then stated that, since conjectures express valuable areas for research, they are important.

    - Some of the most creative minds of this century have been engaged in such research and speculation since Turing established his test, and the case is still being argued.

- INPUT has reviewed many of the arguments pro and con on the subject and makes the following observations:

  - At the time that Turing made his conjecture, the art of computer programming was not very well understood. It was assumed that computers would "do what you told them to do" even if what you told them to do was "something new."

  - There is no question now that machines can surprise us. Unfortunately, most of these surprises have been of the information variety. They are called bugs, and they admittedly are present in programs of even moderate complexity.

  - It is INPUT's intuitive judgment that information surprises will be a major factor in the acceptance of all machines and systems exhibiting artificial intelligence, and especially if such machines are deemed to be thinking. It is likely that humans will be less tolerant of errors on the part of machines than they would on the part of other humans--if for no other reason than that humans are better able to cover their mistakes.

- During the giant "electronic brain" period of the 1950s there was considerable speculation about whether computers could think, and the public and most "experts" concluded that computers can do only what you tell them to do. Thinking-machine research disappeared into the universities under the name of "complex information processing" and "machine-aided cognition." Then John McCarthy coined the term "artificial intelligence" in the 1960s, and it stuck. However, even today there are those who feel the term is inappropriate and might stir up the old public controversy again. It probably will--if it already hasn't.

B. ARTIFICIAL INTELLIGENCE DEFINED

- The best way to define AI is through the research areas that have received emphasis over the past twenty years. However, AI has its roots in the desire of behavioral psychologists to model human cognitive behavior, which involves hypothesizing, reasoning, perceiving, associating, speculating, learning, guessing, and drawing conclusions. These are complex processes that vary significantly from individual to individual, to the degree that measurements can be made. Furthermore, implementation in the "hardware" and "software" of the human brain remains something of a mystery, beyond the left brain and right brain generalizations and understanding of sensory locations. (In other words, we don't have a very good storage map, much less an understanding of how the program works.)

- AI has fundamentally been an effort by humans to create machines in their own image, at least to the extent of having them exhibit human cognitive behavior. Some of the major subfields of AI in which modeling has been done are:

  - Natural language interpretation, in which intelligence is combined with semantic and syntactic understanding to produce quasi-intelligent computer-to-human communication.

  - Image signal interpretation, in which intelligence is combined with visual imaging and pattern/object recognition to permit computer analysis of, and response to, the external world.

  - Speech interpretation, in which intelligence is combined with aural processing and recognition to permit verbal human-to-machine communication.

- Robotics, in which intelligence is combined with mechanical servo-mechanisms to produce "by-reason" rather than "by-rote" machine physical processing.

- Knowledge-based systems, or expert systems, in which intelligence is applied to computer solutions of problems whose complexity or inexactness would normally require human knowledge and reasoning to solve.

- When reviewed against the interests of behavioral psychologists in human cognitive behavior (hypothesizing, reasoning, learning, etc.), it becomes apparent that the intelligence applied in most of these research areas has yet to address the primary issues, except to expose the complexity of the human cognitive processes.

  - Indeed, at the 1984 American Association of Artificial Intelligence (AAAI) convention in Austin, Texas, Marvin Minsky (one of the many "fathers" of AI) specifically pointed out some of the deficiencies of AI research: "AI has been shockingly deficient in developing theories of machine learning. (For example) only a little AI research has gone into the study of machines that learn by analogy. . . . People avoid hard and important problems just as long as they possibly can."

  - Concern about AI research has also been expressed by IBM vice president and chief scientist Lewis M. Branscomb.

    - Branscomb stated (in Shades of Lady Lovelace) that IS professionals are concerned over the excessive claims about AI. In other fields, such excessive claims have resulted in the public having unreasonable expectations about the technology.

    - Expressing his concern about the limited progress of AI research in a number of areas, Branscomb pointed out that only 70 papers had been presented at the 1984 AAAI conference. At a materials science conference of that size, he felt, probably 700 to 1,400 papers would have been presented on the subject.

- However, the major subfields of AI research do have potential for substantial benefit and should not be undervalued even if the tendency to overrate them is already apparent (at least to Mr. Branscomb). INPUT's observations concerning the selected subfields is as follows:

  - Natural language interpretation of high quality is essential if the Turing test is to be passed. (The Turing test essentially states that if a human sitting at a terminal having a dialog cannot determine whether he or she is conversing with another human or with a machine, then the machine can be credited with intelligence. Otherwise the machine cannot be considered intelligent.)

    - Although "quasi-intelligent" conversation remains the current goal of natural language interpretation, natural language systems do not as yet have the ability to pass the test, and there is considerable speculation about whether they ever will, and what it would really prove anyhow.

    - Tests other than Turing's "imitation game" have been suggested. These tests would determine the machine's basic reasoning power by having the machine explain why it reacts or responds the way it does.

  - The fact remains that conversations between humans and machines are at best strained and are generally unnatural. The respected computer scientist Dr. Edsger W. Dijkstra considers computer anthropomorphism (as represented by the desire to converse with computers in natural language) to be unfortunate. He decried the fact that humans take great pains to make computers as unlike machines as possible. Dijkstra remarked that humans disguise the greatest strength of computers by attempting to make them humanlike.

- It is INPUT's opinion that the quest for natural language communication between humans and machines not only is costly but would prove unsatisfactory for both the person and the machine.

- However, the author remembers the first time Autocoder on the IBM 705 typed out the following message on the console typewriter: "NO TAG ERRER, ASSUME TAG ERROR" in the middle of a twenty-minute assembly and continued running. The machine had "guessed" right and saved a big chunk of machine time.

- It makes no difference, then, whether dictionaries, punctuation programs, and passing algorithms are really AI. Even limited intelligence and understanding of natural language is valuable, and research should continue.

- It certainly will be possible for machines to "read," classify, cross-reference, and file documents with increased sophistication, and this is going to be necessary as we proceed toward electronic offices.

- Image signal interpretation can be viewed as still being in its infancy when it comes to playing with various shapes and colors of blocks, but substantial practical applications are already possible based on pattern recognition research. It is possible to recognize hand-written symbols sufficient to incorporate automatic reading and encoding of data from documents. Realizing that some hand-written information has traditionally been unreadable by human beings, it is not beyond the realm of possibility that humans may eventually be required either to write so it is machine-readable or to learn to use a keyboard. It is not important that machines "see" and even remember like human beings. In some cases, they probably will be better, and in others they will never be as good. Once again, continued and extended research is highly desirable, but the anthropomorphism deplored by Dijkstra remains a problem.

- Speech interpretation sufficient for many practical applications is already possible, and full transcription of voice documents is certainly attractive. All of the comments concerning natural language and image signal interpretation apply here, but a few additional observations are required.

  - It may be both practical and useful to use "voice prints" for withdrawing money from the bank or unlocking an automobile.

  - But there are a lot of speech interpretation applications seem to have limited practical value. For example, except in situations where an individual's hands are unavailable (e.g., while driving or flying an airplane, or in some cases of physical handicap) there is little need to activate a machine by voices. In most instances, punching a few buttons would suffice--for example, in dialing a telephone or even looking up a number.

  - Another important issue remains regarding speech interpretation: there is a whole class of people who claim to be waiting to talk to computers before they will use them. It is INPUT's opinion that these are precisely the people who will never know what to say to their computers anyway.

- Robotics in its ultimate form would integrate all of the other sensory and cognitive research into a neat package that would emulate a human being or perhaps improve on the model.

- It should be pointed out here that a great many of the concepts of robotics (of all varieties, from automated machine tools to the more advanced versions) came from Norbert Wiener, who can be called the father of cybernetics (which existed prior to AI).

- He drew the parallel between biological and electrical systems in terms of feedback loops, which are so essential to any complex system. (Wiener remained cautious about the human-machine relationship until his death, especially in supporting decision making, as proposed by his friend John von Neumann.)

- The area of robotics has a broad range of applications ranging from practical factory automation systems, to the whimsical personal servant, to computers that build new (and perhaps better) computers.

- The area of knowledge-based systems, or expert systems, is where the real activity is these days, and the next section will explore this AI subfield in detail. However, there are two important general considerations in such systems:

  - It is easier to "hide the human being in the mechanism" in such systems and pretend the systems are exhibiting intelligence, and you can be sure this will be done.

  - It is more difficult to detect "unpleasant surprises" in an expert system than it is in, say, a robot. And it is also easier to hide mistakes (either intentionally or unintentionally).

  - The combination of these two factors has high potential both for hardware/software/knowledge problems with the artificial system and for human skulduggery. This is a bad combination.

- Before proceeding with a more detailed analysis of expert systems, it is necessary to briefly review some parallel developments that will have impact on each other.

C. A.I., O.R., AND C.C.

- There is a common heritage of artificial intelligence (AI), operating research (OR), and command and control (CC) systems. They all derived from the United Kingdom during World War II, and many prominent persons crossed paths working on different parts of the "messy problems" associated with the more mobile warfare that was developing on land, sea, and in the air.

  - Turing was working on the machines that would be used to break the Germans' Enigma codes.

  - Wiener was working on problems associated with controlling anti-aircraft fire by feedback loops for position prediction.

  - Breaking the codes required new techniques for intercepting both aircraft and U-boats if the real payoff was to be achieved, and the term "operations research" was born.

  - John von Neumann was everywhere, meeting with both Turing and Wiener. He was concerned with the development of computers and with the atomic bomb. Von Neumann was already the leading expert on game theory, which would become increasingly important in CC systems.

- OR became popular after World War II, and its techniques were extended to a number of pragmatic problems in industrial engineering, to various operational problems in the areas of transportation and communications, and to other complex scheduling problems. There has been relatively little contact between OR and the more advanced aspects of AI, but now there is a call from both sides for improved communications between them. There is little question that there is significant convergence between the two general disciplines. However, many areas of AI research are proceeding without the benefit of OR experience, and most OR professionals are not acquainted with AI progress.

- John von Neumann left his mark on the nuclear strategy of the United States. War gaming and strategic planning have been based on game theory concepts. ICBMs leave no time for parliamentary debate (or even for very many telephone calls); computerized decision support systems are inherent in nuclear deterrence. From there, AI has filtered down to CC systems at the battlefield level.

- It is clear that the most complex and crucial expert systems have been developed by DOD. Except for learning of a few wild rumors concerning malfunction and the impact of high-stress war gaming on the human experts playing the games, the general public and business community are unaware of how well such systems work. Indeed, a good argument can be made that such systems will be full of "surprises" if they are even put into operation. However, one thing is certain: there comes a point when a human must trust the system enough to take action or the system becomes an impediment to decision making. It is probable that expert systems will demonstrate many of the attributes of CC systems. That is where decision theory, computer analysis, and AI have been applied.

- In commenting on the attendance at the 1984 AAAI conference, an observer estimated that approximately one-half of the attendees were from colleges or universities; one-third were from government or industry, or were sellers or users of AI (with the largest group of these from DOD); and the remainder were from research institutes. He then went on to say: "Still largely absent, however, are representatives of the professional, industrial, and financial communities toward whom much of this activity is targeted."

IV KNOWLEDGE-BASED, OR EXPERT, SYSTEMS

A. STATUS AND TERMINOLOGY

- There are two major categories of vendors associated with expert systems: those who apply AI techniques to the development of specific application systems, and those who provide the hardware and software tools to assist in the development of such systems. Since the hardware and software tools support research activities and apply to all of the AI subfields (and have been around longer), they have substantially more revenue. The companies that claim expertise in the development of knowledge-based, or expert, systems have normally been spawned out of the universities with many of the principals continuing to hold faculty positions.

- The sudden rush of activity directed toward applying AI to business problems seems to have grown from a number of factors including the Japanese Fifth Generation computer effort; an entrepreneurial spirit; the reported success of a few research projects; and a rush to preempt names such as Cognitive System Inc., Teknowledge, Inc., and IntelliCorp.

  - The "threat" of the Japanese Fifth Generation project provided a focal point of commercial attention for AI, and much of the current activity can be classified as the response of the free enterprise system.

- However, the sudden activity in AI has also been inspired by the successful commercialization of genetic engineering. The winter 1985 issue of AI Magazine contains an article "Comparing Artificial Intelligence and Genetic Engineering: Commercialization Lessons." The abstract remarks that genetic engineering, the academic foundation for the new biotechnology industry, essentially sets an example for AI of the rapid commercialization of an academic subject. There are, otherwise, few precedents for a subject such as AI moving so quickly from academia to the marketplace.

- Then, of course, the media (and market research firms) are also involved in stirring up the pot.

  . The press was out in force at the 1984 AAAI conference. There, John McCarthy, during his President's Address, projected an equation on the screen and remarked that it was "so that the representatives of the press won't think that they've understood everything they've heard (or seen) here." This prompted some later discussions on how press relations could be improved.

  . But, not to worry, the market for LISP machines alone has been forecast by some market research firms to be $1 billion by 1990. Now, with these success figures, what venture capital firm could resist the new terminology? There are some pluses in "spreading the word."

- The computer industry is not unacquainted with the use of new terminology, but AI and expert systems have arrived on the scene with a whole new set. In order to understand expert systems, it is necessary to define some fundamental terms and concepts:

  - Symbolic representations are the substituting of real-world objectives, concepts, and relationships by abstract symbols (letters, numbers, or special characters) that can be manipulated logically by programmed computer processes.

- Symbolic inference is the derivation of new facts or previously undeclared knowledge from existing information through the application of formal and informal logic and search techniques.

- Heuristic search implies the use of rules of thumb for learned and/or experientially derived human expert methods used to narrow a search and guide it toward resolutions.

- Confidence factor is a numeric value associated with a fact or heuristic to indicate an individual or combination weighting of reasonableness or predictability.

- Fuzzy set is a logic formalism expressing the relative set membership of an item in which inclusion or exclusion is not absolute but is represented on a continuum of values.

- Inductive discovery is reasoning from individual observations toward a general principle.

- Deductive reasoning is reasoning from a general principle toward individual observations.

- Public knowledge is that which is generally available through text books, encyclopedia, seminars (proceedings), etc.

- Private knowledge is that which is individually developed, usually through learning, experience, experimentation, etc. (and through what you have probably called "systems analysis" now being referred to as "knowledge engineering").

- During the days when the mystery of the new "electronic brains" was being dispelled, the president of a major transportation company concluded that computers were really no mystery at all. As he so aptly put it, "All a computer can do is say: 'tis, 'tain't; 'tis, 'tain't..." The simplest explanation of an expert system is that it can say "maybe" and then proceed from there. In other words, the human brain is not binary; it does not have a "'tis--'tain't" process, but a "maybe--I wonder" process. The expert system attempts to emulate human decision making and give expert "consultation and advice" in a well-defined and narrow domain of expertise.

- Expert systems have two major components, the knowledge base and the inference engine, but can be supported by other subfields of AI such as natural language interpretation, image signal interpretation, speech interpretation, etc., as shown in Exhibit IV-1.

    - The knowledge base is normally described as being a loose and generally unstructured set of facts, assertions, inferences, observations, hypotheses, rules, and procedures specifically related to a particular domain of activity (such as medical diagnosis). However, it soon becomes necessary to confront the fact that some structure is required, and this has resulted in the following:

        - There is a recognized need for expert systems to have data base access.

        - Both expert systems and data base management systems are recognized (by at least some of those involved in AI) to rely on data or knowledge representation models, and they both involve some inference making.

        - However, they have evolved separately and have placed emphasis upon different aspects of the decision-making process. The result is that it is difficult to integrate existing

EXHIBIT IV-1

AN EXPERT SYSTEM SUPPORTED BY A.I.

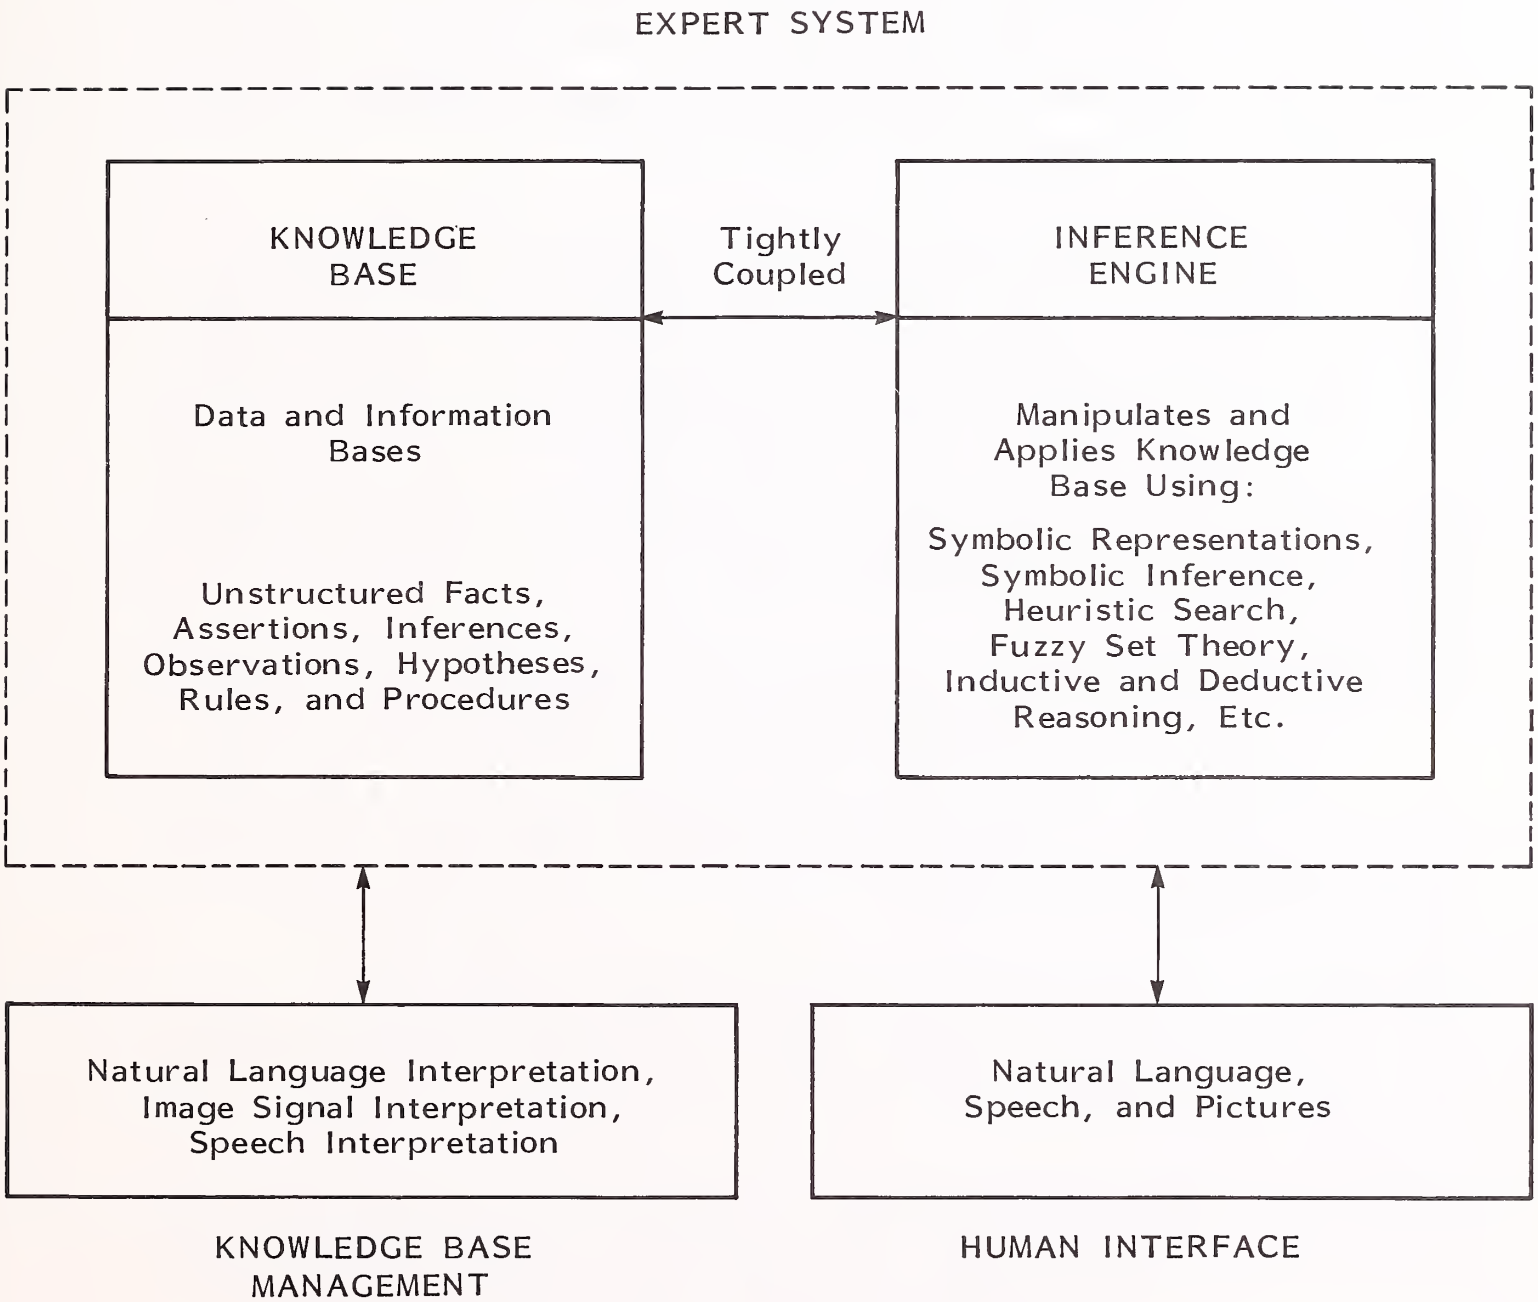

expert systems (of which there are precious few) with existing data base management systems (of which there are many).

- This has resulted and will continue to result in the development of new data base management systems for expert systems.

- The Japanese Fifth Generation effort has opted for the relational model (with at least some hardware implementation), and this seems to be a wise choice (although performance of even the fifth generation may be strained).

- The inference engine manipulates and applies the knowledge bases in progressing toward a problems solution (or decision). It extracts, matches, and applies (tests) the facts, assertions, hypotheses, rules, etc. contained in the knowledge base, against the statement of the problem, refining both the problem statement and the solution as it goes along. Fundamentally, the techniques used are those listed under the definitions of terms and concepts.

- The potential support of expert systems through other AI subfields is obvious when one considers the human interface required for human-machine interaction, and the desirability for computer assistance in screening and structuring of public knowledge sources on specific domains. For example:

  - The human expert should be able to pose a problem in natural language and enter one's own knowledge in the simplest form (pictures, charts, etc.).

  - Scanning, classifying, and reducing data or information, and storing printed material and audio-visual information to build domain knowledge bases (from both public and private knowledge) would do much to facilitate the development of expert systems.

- During the research for this study, a popular business publication listed the transportation company of "'tis--'tain't" fame as being a user of expert systems. A telephone call was made to the company's IS vice president.

  - The call resulted in the respondent's admission that, despite a lot of talk about expert systems and AI (the respondent had recently attended a meeting at which a retired Air Force general had discussed the subject), he was not really sure what made an expert system different.

  - When told that a publication had listed his company as an expert system user, he stated that they must have been referring to a dispatching system they used, but he commented: "As far as I am concerned, that system is straight LP (linear programming)." The company has, we might add, an OR unit under the IS department.

  - The respondent's evaluation of the dispatching system was that it did a better job than some of the newer dispatchers but could not compete with the really experienced ones. Its main value was determined to be in training, where it permitted dispatchers to be moved to new territories and become productive more rapidly.

- There are clear distinctions between the LP dispatching system and the technical definition of an expert system. However, some of the research projects on expert systems have substantially revealed the same result: the expert system can produce results better than X percent of the human experts, but in many cases the Y percent of human experts who are better than the system are expected to handle the harder cases. (In other words, the system refers the case to, for example, a specialist in the medical profession. Or, using the transportation company example, assigns the difficult territory to an experienced dispatcher. In which case, the computer-based system derives its primary benefit as a training or screening device.

- It is obvious that the selection of domains for the development of expert systems is very important. Aside from a few widely publicized systems in medical diagnosis (MYCIN and CADUCEUS) and in mineral exploration (PROSPECTOR), there are currently few new applications. An attendee at the 1984 AAAI conference commented: "The problem of AI seems to be within AI; very few applications of AI tools, successful or not, are reported. I hear some conjecture that the applications in place are so important that their success is being held in strictest confidence. More realistically, experience leads one to expect that when more has been accomplished, more will be reported."

B. HOW EXPERT SYSTEMS ARE DIFFERENT

- Essentially, if you have a problem with a specific solution and can define an algorithm to solve it, you do not need an expert system. Expert systems are for "fuzzy," complex problems in which human judgment and reasoning are normally employed (with varying degrees of success). The algorithmic nature of OR problem solving would therefore rule out OR solutions as expert systems.

- In addition, there is a set of general characteristics that distinguish expert systems from conventional data processing (DP) systems that may solve rather complex problems. These characteristics are highlighted in Exhibit IV-2.

  - Expert systems display expert performance by solving problems that would normally require a human expert. Unlike a traditional DP application, which is rigidly defined and leaves no room for probabilities or uncertainties, expert systems arrive at an acceptable solution, in a reasonable period of time, despite ill-defined data and processes.

EXHIBIT IV-2

DISTINGUISHING CHARACTERISTICS OF EXPERT SYSTEMS

- Display Expert Performance
- Employ Symbolic Representation
- Employ Symbolic Inference
- Use Heuristic Search
- Have Extensive Domain Knowledge Foundation
- Fall Back to "Weak Reasoning" for Answers
- Reformulate Problems
- Have Self-Knowledge

- Conventional data processing systems use arithmetic symbols and processes to represent data elements, files, and logic paths. Whether mathematical algebra, Boolean algebra, or relational algebra is used, they do not cope very well with abstract and vague concepts and relationships. Expert systems, however, employ symbolic structures to represent both concrete and vague abstract concepts, hypotheses, associations, and so on.

- Data processing normally employs algorithmic processing, which means that the process (including handling of errors) is predetermined. Indeed, when surprises occur the program frequently aborts. Expert systems employ nondeterministic symbolic inference, which means that the sequencing of operations and the logic applied is dynamically determined by the system. Surprises are expected as a result of partial problem solutions and external events; the system adjusts itself and continues toward a solution.

- Conventional data processing systems normally employ a single-path approach, or an exhaustive search of all possible candidates in seeking a problem solution. Expert systems use heuristics or "rules of thumb" to narrow the search and guide it toward a workable--though perhaps imperfect--solution. Because of heuristic searching, expert systems can address large, more complex, fuzzy, and nondeterministic problems that could not be practically addressed by the exhaustive searches of traditional DP approaches.

- Expert systems have extensive domain knowledge, both public and private as their foundation. (This actually represents a shift from an emphasis upon laws of reasoning and inference, which resulted in the development of earlier cognitive systems, and is comparable to the early DP experience of developing general purpose management information systems without underlying data.) DP systems usually only concentrate on specific data required for actual processing, even if those data are the result of an extensive data requirement study.

- When heuristics associated with a domain fail to result in a successful search, expert systems normally fall back to "weak reasoning," based on general logic, to provide some answer. Conventional data processing systems do not provide answers if these exhaustive searches fail. This is possible because expert systems reformulate the human problem representation into one suitable for more flexible computer resolution.

- Expert systems have knowledge of how they have resolved problems, and they explain what they have done and justify the recommended solutions. This self-awareness is not required in algorithmic systems, because these systems have been "told what to do."

• Expert systems are usually applied to problems that also exhibit a set of characteristics that set them apart from those usually identified with algorithmic processing systems.

  - Problems addressed by expert systems are usually large and complex, with multiple and possibly conflicting solutions.

  - The problems are not predictable, being fuzzy and uncertain, and these solutions require symbolic inference and abstract logic.

  - The problems are data intensive, requiring large amounts of factual and experiential data and associated logic representation.

  - Because of size and complexity, the problems normally can be resolved only by using heuristic search; they require "expert" judgment and are not amenable to algorithmic solution.

• At this point, it seems appropriate to point out that as more data become available for both planning and control purposes, practically all problems become complex and uncertain. As personal computers are used to create distributed and personal data bases, these may be viewed as knowledge bases in the sense that they represent a particular "way of looking at things"--a way that an individual (or organization) uses to make decisions. In other words, as more data and information become available in the decision-making process, problems become fuzzy and more alternatives become available. It is obvious that the purported merits of expert systems are as logical extensions of decision support systems. The question becomes: Are they a practical extension?

C. HARDWARE AND SOFTWARE TOOLS

1. PROBLEMS OF IMPLEMENTATION

- By definition, expert systems are complex, and it should be obvious that if the problems they are attempting to solve were easy more progress would have been made. The components of an expert system are best depicted with a general flow diagram, as illustrated in Exhibit IV-3. In its ideal implementation, such a system goes well beyond all of the advanced work that has been done in fourth-generation languages, data base management systems, and operating systems (the scheduling is substantially more complex). Essentially:

  - The language processor provides a narrow problem-oriented language specific to the expert user.

  - The "blackboard" provides for the recording of intermediate results.

  - The knowledge base contains facts, heuristic planning, and problem-solving rules.

  - The interpreter applies the rules.

EXHIBIT IV-3

THE COMPONENTS OF AN EXPERT SYSTEM

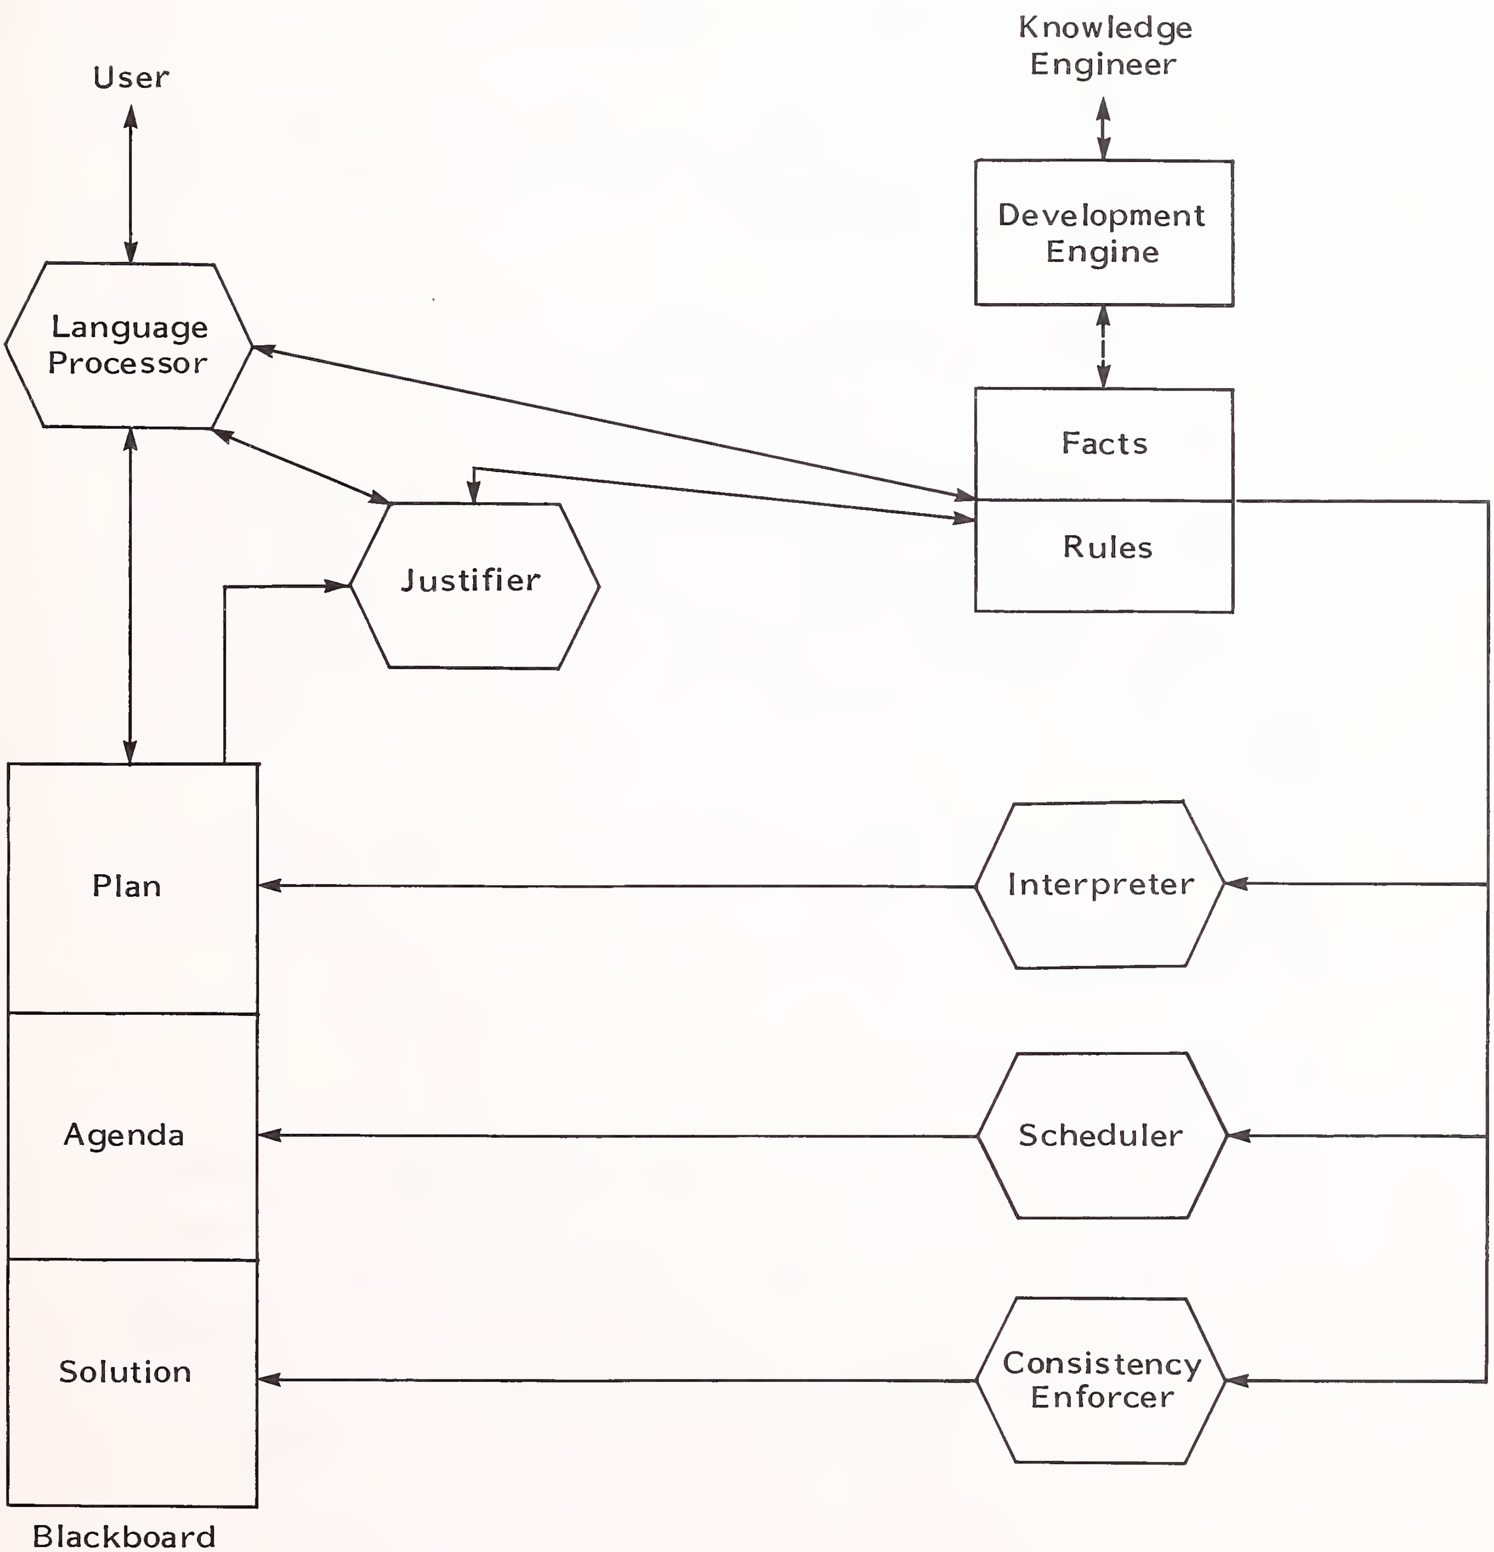

- The scheduler controls the order of rule processing.

- The consistency enforcer adjusts previous conclusions when the knowledge (or data) base changes.

- The justifier rationalizes and explains the systems behavior.

- Considering the fact that each domain (application) is highly specialized, it is understandable why the number of expert systems developed to this point has not been large.

- The reason that expert systems are receiving so much attention is that hardware and software tools to facilitate development are beginning to appear in the marketplace. Before discussing these tools it should be pointed out that there is currently an extreme shortage of people who are knowledgeable about these tools or who have any ideas of how to develop expert systems. Indeed, the tools themselves add another level of complexity in the development of such systems. (The situation is comparable to having MVS/XA, IMS and PL/I descend upon you at one time without any systems programmers being available.)

- Therefore, even though expert systems may be the logical extension of decision support systems, available tools and personnel will determine whether they are practical.

2. LISP AND LISP MACHINES

- When people hear AI mentioned, most of them think of LISP. As one observer at the 1984 AAAI Conference stated: "It may not be entirely correct that AI is what LISP makes feasible, (but) the coincidence of the capabilities of LISP with accessibility to large-scale computing power on very economical terms has provided much of the energy for the AI explosion."

- LISP is not a new language. It has been around nearly as long as FORTRAN. Unfortunately (or fortunately, depending upon how you want to look at it), LISP has not been formalized. It has been left to grow according to local requirements, most of which exist in the universities and research centers. There are three important characteristics of LISP:

  - The ability to evaluate words and lists, rather than numbers and tables.

  - A facility to handle truth tables and Boolean algebra (the "if--then" operational concept which most people are familiar).

  - An easygoing environment that makes it easy to stop, take stock, and make changes while proceeding (during development).

- Of course, LISP isn't the only symbolic language (in fact, LISP itself includes many languages), and the Japanese have elected to support PROLOG in their Fifth Generation project. At the "first national meeting" of the Symbolics LISP Users' Group (SLUG) a progress report stated the following:

  - "High-performance" PROLOG would not be out soon for the Symbolics 3600. (Symbolics is only one of several companies currently marketing symbolic processors.)

  - PROLOG KR (which was developed in Japan), however, was immediately available.

  - In addition, LOGLISP, a combination of LISP and PROLOG developed at Syracuse University, was reportedly available.

  - And an ADA project was reported to be well underway.

- For anyone having fought battles over language standardization, it is readily apparent that neither LISP nor other languages for symbolic processors have been formalized. The hope has been expressed that the commercialization will lead to formalization, but right now the users seem to be making demands for a polyglot, and the hardware vendors may be hard put to deny their wishes.

- Exhibit IV-4 lists some of the workstation hardware currently being marketed to support the expert systems development market. Personal computers microcoded with LISP and PROLOG are currently under development, and this should bring the cost down substantially.

- However, LISP has addressed primarily the interpreter portion of the expert system development problem; and, as mentioned previously, the building of the knowledge base is now recognized as all important. Fortunately, it was recognized that the "shell" of hand-tailored expert systems could be salvaged, and that has led to the emerging expert systems software industry.

3. EXPERT SYSTEMS SOFTWARE

- The importance of knowledge in intelligent problem solving (using expert systems) is summarized in Exhibit IV-5. Acquiring and encoding knowledge (by those scarce knowledge engineers) are the most difficult and costly steps in building expert systems--just as problem analysis and data base development have been in conventional data processing systems.

- This has led to the development of a number of commercially available software shells, which are beginning to appear in the marketplace. Some of these tools are listed in Exhibit IV-6. Although a detailed analysis is beyond the scope of this study, some basic analysis is warranted; it is obvious from their prices that they vary substantially in capability and hardware supported. However, they can generally be categorized as follows:

  - The micro-based systems (Expert Ease, M.I. Personal Consultant, and TIMM) typically process knowledge represented as IF:THEN and use backward chaining. Essentially, the system starts with a hypothesis,

EXHIBIT IV-4

EXPERT SYSTEM HARDWARE

VENDOR	EQUIPMENT	PRICE
LISP Machine, Inc.	Lambda Lambda 4 X 4/+	$66,500 $45,000
Symbolics, Inc.	3600	$59,000 - $84,000
Xerox, Inc.	1108 (Dandelion) 1132 (Dorado) 1100 (Dolphin)	$25,000 - $50,000 $150,000
DEC	VAX	
Tektronix, Inc.	4404 Artificial Intelligence System	$14,950
Texas Instruments	Explorer	$52,000 and up
Perq Systems Corp.	Perq AI Workstation	$40,000

EXHIBIT IV-5

IMPORTANCE OF INTELLIGENT PROBLEM SOLVING

- Knowledge = Facts plus beliefs plus heuristics

- Success = Finding a practical solution - one that is good enough and is not too expensive

- Search efficiency directly affects success

- Aids to Efficiency:

  - High Quality Knowledge Base
  - Elimination of Blind Alleys and Redundant Computation
  - Improved Speed of Operations
  - Coordination of Knowledge Sources
  - Various Levels of Reasoning

- Contributor to Problem Difficulties

  - Bad Data or Knowledge
  - Changing Data
  - Alternatives that Must Be Evaluated
  - Procedural Complexity

EXHIBIT IV-6

EXPERT SYSTEMS SOFTWARE TOOLS

VENDOR	PRODUCT	HARDWARE	PRICE
Expert Systems, Inc.	Expert Ease	IBM PC, DEC Rainbow, Victor Technologies 9000	$125 $2,000
Teknowledge	M.1 S.1	IBM PC Xerox 1100 & 1108 and DEC VAX	$12,500 $50,000 - $80,000
IntelliCorp	Knowledge Engineering Environment (KEE)	Symbolics 3600 Xerox 1100 Xerox 1108	$60,000
Texas Instruments	Personal Consultant	TI Professional Computer	$3,000
Inference Corp.	Advanced Reasoning Tool (ART)	Symbolics 3600 LISP Machine DEC VAX	$48,000 - $60,000
General Research	The Intelligent Machine Model (TIMM)	DEC VAX IBM PC	$39,500 $9,500

finds a rule whose premise (IF) supports the hypothesis, and then attempts to verify the premise by searching the knowledge base for a fact that supports the premise. If the fact isn't found, the system searches for a rule that can be used to infer the fact. The process continues until the hypothesis is either verified or disproved. Such systems provide only limited communication with the knowledge engineer concerning flow control (the order in which the system asks questions).

- S.1, obviously a more expensive system for more powerful processors, provides for interpretation of control terminology similar to that found in regular programming languages.

- Backward chaining works well for problems where selection from a number of possibilities is appropriate for the domain (applications). However, sometimes the application cannot be specific in that manner, or the number of possible solutions is too large to handle on a micro because of processing or space limitations. In such cases, the application is appropriate only for developing prototype systems or for training knowledge engineers.

- The more advanced commercial shells such as KEE, ART, and S.1 incorporate multiple inferencing techniques by supporting not only backward but also forward reasoning. This means that the user may start with a fact and find a rule whose premise is verified by the fact. Thus, the more advanced models are appropriate for more applications.

- The software development shells are attached to the knowledge base being built. (Refer back to Exhibit IV-3.) Once the expert system (application) is completed, the development engine is removed and the system is turned over to the users. The question now becomes: What type of applications are we talking about?

D. APPLICATION AREAS

- Ten major categories have been identified for expert systems. These include:

  - Interpretive systems, which infer situations from observables and can be used for many of the expert systems support functions (speech understanding, image analysis, signal interpretation, etc.) as well as a variety of intelligence operations.

  - Prediction systems, which infer likely consequences from current observations and historical trends. The classic weather forecasting problem falls in this area, along with crop estimation and battlefield movement prediction.

  - Diagnosis systems, which infer malfunctions from observables and include not only the classic medical diagnosis but also electromechanical and software failure analysis.

  - Design systems, which configure objects or processes under goal constraints and can be used for computer configuring, financial budgeting, and circuit layout.

  - Planning systems, which explore possible future actions to produce the necessary steps to accomplish a goal. Applications would include automatic programming, project planning, and remote planning.

  - Monitoring systems, which compare observations against normalized criteria to determine deviations. Examples are financial management monitoring, air traffic control, and nuclear power plant monitoring.

- Debugging systems, which prescribe specific remedies for malfunctions, and are especially appropriate for computer software programs and digital circuitry. (If you have tried to determine how and why some programmers can debug so much better than others, you will understand some of the complexity of building an appropriate knowledge base for expert debugging systems.)

- Repair systems, which develop and execute plans of repair for a system. Examples are computer hardware systems and communication networks.

- Instruction systems, which diagnose and "debug" student behavior and enable customized computer-aided education.

- Control systems, which adaptively govern the overall behavior of a system. These apply to industrial production, nuclear power plant control, investment portfolio management, and so on.

• The potential application areas lead to the conclusion that practically all information systems components (hardware, communications, software and paper-based information flow in support of the decision-making process) can be subject to replacement with expert systems, and with impact on the human experts (once they have contributed their knowledge), as shown in Exhibit IV-7.

• The impact on IS personnel is understandable--after all, knowledge-based systems are the logical extension of data base and information systems. However, the heavy concentration of potential impact among financial personnel might pose some funding problems. Of course, our example is exaggerated and arguments will be made that expert systems will merely permit highly skilled personnel to do a better job. However, the threat is there, especially for information systems personnel who choose to ignore the technological advances that are possible in the systems development process.

EXHIBIT IV-7

CATEGORIES OF EXPERT SYSTEMS AND REPRESENTATIVE IMPACTS

CATEGORY	SOME POSSIBLE IMPACTED PERSONNEL
Interpretive Systems	Data Entry Personnel Business Analysts
Prediction Systems	Forecasters (Including Consultants)
Diagnosis Systems	Doctors Field Engineers
Design Systems	Systems and VLSI Designers Accounting and Financial Systems Personnel
Planning Systems	Programmers Schedulers
Monitoring Systems	Project and Financial Control Personnel
Debugging Systems	Programmers Circuit Designers
Repair Systems	Field Engineers
Instruction Systems	Teachers Training Personnel
Control Systems	Industrial Engineers Investment Analysts

- INPUT doesn't think Lady Lovelace would have viewed expert systems much differently than she did the analytical engine. She intuitively knew that programmers would be required, and, likewise, one can almost hear her say: "My, expert systems will certainly require a lot of people to advise them on what they may do! I think I will call them 'knowledge engineers.'"

V INPUT'S ANALYSIS

A. WHERE THINGS REALLY STAND

I. A SIMPLE CASE STUDY

- Perhaps the best example of a practical AI application that has been developed in a university environment and actually transferred to an industrial environment is XCON. XCON is an expert system that expands and validates orders for DEC computer systems and then configures the placement of components within equipment cabinets. The original version of XCON, called R1, was developed at Carnegie-Mellon University in 1979 and delivered to DEC in 1980. A brief description of this project is contained in "R1 And Beyond: AI Technology Transfer At DEC," AI Magazine, Winter 1985.

- The article makes several important points that relate to expert systems emerging from the research laboratories.

  - The system as delivered to DEC in 1980 contained approximately 750 rules. Although it could configure many of the orders, it required "extension and refinement" before it became dependable enough for use.

  - A group was formed at DEC to assume responsibility for maintaining XCON. This group was composed of DEC engineers who were not selected for their AI expertise (although some were at least familiar with AI). These engineers, trained in expert systems programming and knowledge acquisition, corrected and extended the system's knowledge. After about a year, and with help from Carnegie-Mellon, the group was able to take over maintenance of XCON (which by that time had grown to contain approximately 1,000 rules). This transfer was described as successful but difficult.

- The system has now grown to 2,500 rules, and even though it was a difficult experience an industrial user has been able to help complement and maintain an AI system.

- Since 1982, DEC has expanded its AI activities and has developed a number of expert systems including XSEL, which is designed to aid salespeople in tailoring systems to meet customers' particular needs. The experience with XSEL paralleled that of XCON. It was developed at Carnegie-Mellon, delivered to DEC, and established by a highly skilled maintenance group.

- Based on the XCON-XSEL experience, DEC has established a formal training program for knowledge engineers and a research group (Knowledge Engineering Advanced Development Group) to ease DEC's reliance on academic researchers.

- Expectations about AI's capability have increased as people have begun to understand the potential of AI, and as the technology has been employed. In addition, the way people feel about the influence of AI has changed.

- The author of the article, Stephen Polit, warns about exaggerating the capabilities of AI--in order to avoid unrealistic expectations--and urges IS managers to reevaluate goals. He reminds readers that even new ideas and technological developments cannot solve every problem. AI is really affecting the business practices of DEC and many other organizations.

- INPUT has made several notable observations concerning the XCON case study, none of which are intended to diminish the significance of either the achievement or the lessons learned.

  - XCON is essentially the replacement of a group of technical editors who represented a bottleneck in the order entry process. Evidently DEC had unsuccessfully attempted to automate this function before turning it over to Carnegie-Mellon as a research project. (DEC has had a long, historical relationship with AI.)

  - The fact that Carnegie-Mellon was able to develop an expert system to solve the problem does not mean that it is either necessary or even desirable to have an expert system perform configuration control. IBM, for example, exercises configuration control in the order entry process on a more complex product line, with higher volumes, without resorting to expert systems.

  - It is not at all certain that the resulting expert system is a cost-effective solution to what appears to be essentially a data processing problem. And it is probable that no cost analysis has ever been made of the development costs of the system. At this point, such a cost analysis may not even be possible.

  - However, for anyone acquainted with the problems of life-cycle costs, the heavy emphasis upon maintenance associated with both the XCON and XSEL projects is indeed ominous. To maintain the knowledge base, the organization needs highly skilled technical personnel having a rare or nonexistent specialty (AI). This is an extremely important lesson for anyone who would benefit from the DEC experience.

- It appears that knowledge engineers are destined to learn what IS management already knows: maintenance represents over 60% of total systems cost over the system's life cycle. The question becomes one of whether the cost of knowledge base maintenance will compound an already serious problem, and, indeed, whether these complex expert systems will even be maintainable as they grow. The XCON case study points out a specific weakness of expert systems--a weakness that is already recognized in the industry. It is what INPUT prefers to describe as "systems fragility."

2. EXPERT SYSTEMS FRAGILITY

- Current expert systems are frequently referred to as being "brittle," because they tend to fail on problems that are not well centered within their specific domains. This is in contrast to robust solutions that seem to function well at, or even beyond, the periphery of their intended use. For example, queueing networks, which were developed as a tool of operations research and industrial engineering, have been found to yield good results in complex scheduling problems found in computer operating systems--much to the surprise of computer scientists.

- INPUT prefers the term "fragility" as a reminder that expert systems are not robust. Unfortunately, they are not obviously fragile. There is a tendency to package expert systems, which will result in a certain amount of protection for the contents but will also obscure the fact that they may break without warning and be worthless. It may also result in hiding the fact that it has become worthless. The least one can do is stamp the contents: "Fragile--Handle With Care."

- This general labeling of expert systems as fragile is not capricious. Responsible researchers in the AI field are candid about the current state of the art, and at this point the development of expert systems remains "art" despite the knowledge engineering label. Let us consider some of the acknowledged problems associated with current expert systems.

- The current problems associated with brittleness center around the inability of expert systems to reason about either their own knowledge or that of others. They are unaware of what they know or, more importantly, what they do not know. They cannot modify their behavior based upon what the user already knows about the subject. (Imagine dealing with a human expert with such characteristics. This may, of course, be oversimplifying the situation.)

- These defects lead to narrow and restrictive domains for expert systems because they lack flexibility and adaptability. When confronted with problems on the periphery of its knowledge base, an expert system is likely to wander off aimlessly searching for a solution because it "doesn't know what it doesn't know"; or, as an alternative, may complete its processing ungracefully when it cannot apply any of its solutions to the specific problem.

- As a result of narrow and restrictive domains, there is a tendency to develop individual expert systems in which there are substantial overlaps in the knowledge base. For example, there is certainly substantial knowledge base overlap in the XCON-XSEL case study, and yet it was determined that two separate groups of knowledge engineers were required--one for the configuration control system and one for the sales support configuration system.

- Even with expert systems being limited to narrow domains, the human-machine dialogues are so formalized and restricted that a knowledge engineer is frequently required to interface between the human expert and the expert system in order to capture and represent knowledge. Since a knowledge-based system is by definition a feedback-loop-oriented system, it is probable that the development engine will remain a part of the expert system long after it is in productive use and perhaps over the system's life cycle. In the simplest possible terms, the current defects of expert systems indicate a high probability of heavy maintenance by highly skilled personnel.

- The basic problem is that expert systems don't learn very well, and learning is one of the "hard problems" that Marvin Minsky chastised the attendees at the 1984 AAAI conference for not addressing. Research in human learning has failed to reveal how it takes place in the brain. It hasn't begun to explain the mysteries of the human mind. It is INPUT's opinion that current expert systems resemble libraries (complete with human librarians, in the form of knowledge engineers) much more closely than they do the human learning process.

- A step beyond learning is decision making, in which expert systems, despite their artificial intelligence label, adopt what is essentially an algorithmic approach. It is only necessary to examine knowledge representation in expert systems to determine that this is so; and to the degree that expert systems are to be used for decision making, it is important to understand the limitations of algorithms approaches to decision making in the general business environment. Knowledge representation in expert systems is currently one of the primary research areas. A brief summary of the various representation schemes follows:

  - Predicate calculus evaluates arguments and their relationships to a truth table (either TRUE or FALSE).

  - Rules have the classic "IF:THEN" scheme of representing knowledge for purposes of inference.

  - A semantic network is a knowledge representation for describing properties and relationships through the modeling of objects as graph nodes and of their relationships as labeled arcs connecting the nodes. (It is not clear that the human brain works in this fashion, although there seems little question that an analog representation is more likely than the TRUE-FALSE, IF:THEN schemes.)

- A frame is a knowledge representation in which descriptive, relational, and procedural information about an object (or a concept) is grouped into a data structure consisting of various "slots" in which a pyramidal inheritance is inferred. (Although this may seem unclear, it has improved since the term was introduced by Minsky in 1975; later research concluded "it is not at all clear now what frames are, or were even intended to be.")

- A procedural is a knowledge representation in which the execution of a procedure (usually algorithmic) permits the inference of some fact, concept, hypothesis, etc. (In other words, if a procedure sorts on a certain key it implies knowledge concerning the relationship of the key elements.)

- In order to understand the fragility of expert systems, it is important to at least identify major areas of ignorance that contribute to this fragility. In other words, even if expert systems are unaware of what they don't know, it is desirable to understand what the developers (or the proponents) of such systems don't know.

3. IGNORANCE AND EXPERT SYSTEMS

- The areas of ignorance associated with expert systems are so vast and complex that they can only be touched upon in this report. We have isolated three general areas that should not be ignored in considering the practical applications of expert systems. These are learning, decision making, and computer hardware performance. An appropriate single source for identifying these areas is The Encyclopedia of Ignorance, Pergamon Press Ltd., 1978. The importance of understanding ignorance is succinctly presented in an essay on evolution which quotes from The Unwritten Comedy, by Phenella.

- "To be ignorant of many things is expected."

- "To know you are ignorant of many things is the beginning of wisdom."

- "To know a category of things of which you are ignorant is the beginning of learning."

- "To know the details of that category of things of which you were ignorant is to no longer be ignorant."

• Although the above was written to describe a human being, it could as well apply to an expert system, and perhaps provide a rough measure of progress toward knowledgeable systems. The Encyclopedia of Ignorance represents the efforts of some of the world's leading scientists, including Nobel laureates, to categorize areas of ignorance in such diverse fields as space, mathematics, physics, and biology. The authors identify specific categories of ignorance that are important when considering the development of expert systems. These are as follows:

- In the general area of learning, it can be stated that:

    . Learning is dependent upon memory, and we do not understand how memories are stored in the brain. In fact, the more research that has been done, the more complex the problem seems to become.

    . We do not currently understand consciousness (or how the brain and mind relate) or whether it is of any use to us; specifically, we are ignorant of its potential role in intelligent computer systems. We do not know if 1) such systems would benefit from consciousness, 2) if they could be provided with it, and 3) if we humans would recognize it if they had it.

- Despite progress, including that associated with artificial intelligence, we are still ignorant of the languages of the brain. We are aware, however, that it is improbable that brain languages consist of symbols occurring in a single serial string, and that is what symbolic processing is all about.

- John von Neumann was especially interested in decision making in economics and, along with Morgenstern, published Theory of Games and Economic Behavior in 1944. Some of the fundamental problems identified with the application of mathematics to the social sciences were identified at that time and, despite progress (the last Nobel Prize in Economics was awarded to a mathematician), modeling in the social sciences remains an area of ignorance.

  - In 1943, von Neumann stated that economic theory uses mathematics in an exaggerated way, perhaps, and also not very successfully. The only change today has been the provision of quantitative data, and von Neumann's remark has had little effect.

  - In observing that social phenomena are at least as complex as those of physics, von Neumann concluded that mathematical discoveries comparable in magnitude to calculus may be necessary to progress in the field.

  - There remains a great area of confusion (ignorance) in distinguishing mathematical methods appropriate for modeling physical (mechanical) phenomena and their incorporation in systems using game theory and/or the social sciences.

- Essentially, the social sciences deal with artificial systems--those involving interaction with human beings--and these are precisely what expert systems are all about.

- It can be prove mathematically that artificial systems are unpredictable in their total behavior, especially in regard to the actions that human beings will take.

- In addition, it is impossible to determine the state of the total system, even by applying probabilistic judgments as in game theory.

- Hardware performance of even conventional data processing systems has been of concern to INPUT, and we have frequently quoted Hans Bremermann's article, "Complexity and Transcomputability." However, the points made bear repeating, since they apply to expert systems.

  - Transcomputability is described as the point beyond which the computational cost exceeds all bounds that govern the physical implementation of the algorithm.

  - It can be shown that many of the algorithms of artificial intelligence and operations research are transcomputable--in fact, any algorithm whose computational cost grows exponentially with a size parameter n is transcomputational for all but the first few integers of n.

  - Transcomputability applies not only to existing computers but to any that can even be built (including the Fifth Generation system).

  - Bremerman concludes that it is disturbing to imagine that transcomputability applies to all computers and that many people chose to ignore this fact.

- In summary, knowledge-based systems are being developed, while we are still ignorant in many important areas that seem necessary for correcting the systems' currrent weaknesses (flexibility, adaptability, cost, and even practicality). It is important to approach expert systems with understanding of these weaknesses.

4. DECISION THEORY AND GAME THEORY CONTROVERSY

- It is well beyond the scope of this study to analyze decision and game theory in any detail, but expert systems, as extensions of decision support systems, will have a propensity to head in that direction.

- A few paraphrased quotes will help to identify certain concerns that INPUT shares with many of those most closely associated with the continuing controversy on the subject.

- In 1955 John von Neumann warned that the best computers could expect to do in the foreseeable future was to provide mechanical aids for decision making while the process would of necessity remain human. He pointed out that no automatic approximation existed for the human intellect, and that it was impossible to describe the logic involved in "intuitive" decision making.

  - In a 1980 biography of von Neumann (John von Neumann and Norbert Wiener), Steve J. Heims pointed out that von Neumann had thus limited theory by the inability to develop computer models of the human brain. Heims then noted that others had not been as cautious in applying game theory to human decision making.

  - Gregory Bateson, who shared Norbert Wiener's concern about the application of game theory (especially to military strategy), felt that the mere use of game theory rules and premises would change the "players'" perceptions of how they dealt with each other, and that they would become unable to relate on a more human basis. He anticipated that in the long term such changes would be in a "paranoidal direction and odious."

- Heims then reaches the conclusion that most large organizations (corporations and governments) are vulnerable to seeking simple, automatic solutions to problems that require decisions--that there is a natural tendency to diffuse responsibility. He then warns that miscalculation of the effectiveness of pseudoscientific solutions to complex problems is extremely dangerous.

- As INPUT concluded in Impact of Office Systems on Productivity (1983): "In summary, we conclude that significant improvement in white-collar productivity depends upon a complex hierarchy of interrelated systems, all of which are approaching current limitations of technical fields. Information systems organizations are going to have to deal with questions of artificial intelligence whether or not the questions can be answered or the problems solved." Good Luck!

B. RECOMMENDATIONS

1. WHAT IS GOING TO HAPPEN?

- Recently, a local Silicon Valley newspaper--the Times Tribune, Palo Alto (CA) featured an article entitled "Investors Losing Patience in the AI Software Market" (Feb. 25, 1985). So perhaps this section should carry the subhead "What Is Happening." The article pointed out the following:

  - Few AI-oriented software companies started in the past few years have made money, and most have not produced marketable products.

- Continuing needs for cash are meeting resistance from the venture capital community.

- It is shakeout time already, and AI itself seems to be getting a bad name (before anybody really understands the products very well).

- There is a trend away from the use of the term "AI" in favor of such vague terms as "smart" and "enhanced."

- The reasons given for the problems of AI software companies are not unusual: technical problems in development and conflicting ideas about the market.

• INPUT does not feel that AI--or, specifically, expert systems--will be forced back into the universities and research centers for another 20 or 30 years because of lack of funding. However, Lady Lovelace was right. We do tend to overrate what we find interesting or remarkable and then to undervalue the device's true condition. This report notwithstanding, INPUT believes we are still in the overrating stage of the process as far as expert systems are concerned and that this will continue until a number of significant systems failures (as opposed to vendor failures) occur. During the overrating period:

  - A lot of tools for the development of expert systems will be sold (LISP machines and development shells). Gradually the sale of such tools will go beyond the universities and research centers and find their way to major software development companies and even enterprising end users.

  - The software companies that attempt to develop general purpose expert systems will continue to experience both technical and marketing problems.

- A number of major expert systems efforts will fail--some with a bang and some with a whimper.

- Expert systems will be developed for problems that, more than 95% of the time, could be solved with other approaches (operations research, data base systems, systems analysis, etc.) more effectively and economically.

- Expert systems that do become operational will require excessive resource use in terms of machine time and personnel costs for maintenance.

- Expert systems (or reasonable facsimiles thereof) will be blamed for failures in the decision-making process regardless of whose fault it was.

- The IS department (or knowledge-engineering department) that developed (or installed) the expert systems will be blamed for the failure.

- Expert systems will then become undervalued by everyone involved, and another Lovelace cycle will have been completed.

- If all of this seems familiar, it should. The same cycle exists for other magic solutions to complex problems.

2. WHAT YOU SHOULD NOT DO

- Do not ride the "overrate-undervalue" rollercoaster this time. It is becoming more dangerous as the "solutions" involve higher levels of management in artificial systems (having the potential weaknesses INPUT has described).

- Do not remain ignorant of what expert systems can do--either for you or to you.

3. WHAT YOU SHOULD DO

- This report has identified several categories of ignorance that underlie systems. This is the first step in learning. Managers are urged to become familiar with the details of the categories of things they are ignorant of--that is, to no longer be ignorant. Essentially, this is to understand the fragility of expert systems so they can be handled with care. However, this will also make the manager an expert on expert systems. That is the paradox of understanding ignorance.

- It is important to understand the fragility of expert systems, and this report has emphasized weaknesses to take the peak off the overrate phase of the Lovelace cycle. But it is also important to obtain knowledge about expert systems and their use. It is recommended that responsibility for acquiring this knowledge be centralized in a separate research function (either a single person or an entire organization) with responsibility for the following areas:

  - Operations research and statistics.

  - Information theory.

  - Decision theory.

  - Library science.

  - Artificial intelligence (including expert systems).

- If there is already an operations research group in your organization, it should be organizationally located (or relocated) within the IS department and assigned responsibility for the other areas. Staffing of the other areas will depend upon the particular requirements of the company. However, it is not deemed essential to rush out and recruit knowledge engineers (especially since they aren't available). The purpose of this group is to become knowledgeable in order to enhance and extend information systems to more complex problems--not to sell expert systems.

- The organization should be viewed as an applied research group to work with data base administration, information centers, development centers, and libraries, as well as with IS and user development efforts. Ideally the funding for research in advanced techniques would come from consulting fees (or charge backs) from other areas. (DEC had problems with this because of high demand for knowledge engineers and had to isolate a separate AI research function, but the broader scope of this recommendation should alleviate some of the problem.)

- It would seem that information and development centers are the ideal vehicles for promoting the use of advanced techniques, for identifying promising applications areas, and for exercising quality control of systems being developed in a distributed systems development (DSD) environment.

  - Initial expert systems available on a commercial basis will be addressing the financial and planning functions. These are precisely the types of applications that are addressed by information centers (and the DSD environment in general). Knowledge engineers (systems analysts) from the research organization should begin to explore how the information (reports) from the DSD environment is being used to support the decision-making process, and whether expert systems (or other advanced analysis tools) can be applied. Then it can be determined whether commercially available AI products are of value.

  - The tools of AI (LISP machines and software shells) are being promoted for productivity improvement. The research group should evaluate whether such tools are appropriate for either information or development center use.

- New Opportunities for Software Productivity Improvement details the quality control problems associated with the DSD environment. All of these problems will be accentuated with the use of knowledge-based systems. For example, one of the problems isolated in the DSD environment is "conflicting reports to management." Imagine the problems that would arise if "conflicting decisions (solutions)" were to become available from overlapping expert systems.

- In fact, the strongest recommendation involves quality control of vital data, information, and knowledge. Readers are urged to take a look at the INPUT report New Opportunities for Software Productivity Improvement and to make an applied research function an extension of systems quality control; many of the advanced tools and techniques are required for quality assurance.

- A good argument can be made that as more data, information, and knowledge become immediately available at the desk top, the choices (decisions) necessary will overwhelm the human component of artificial systems. Expert systems may be the only answer at the point of human-machine interface that will narrow the choices; and the development, installation, and management of such systems is probably going to be your responsibility. The last advice is quite simple: Handle expert systems with care--not only because the contents may be fragile, but because the package may contain a bomb!

About INPUT

INPUT provides planning information, analysis, and recommendations to managers and executives in the information processing industries. Through market research, technology forecasting, and competitive analysis, INPUT supports client management in making informed decisions. Continuing services are provided to users and vendors of computers, communications, and office products and services.

The company carries out continuous and in-depth research. Working closely with clients on important issues, INPUT's staff members analyze and interpret the research data, then develop recommendations and innovative ideas to meet clients' needs.

Clients receive reports, presentations, access to data on which analyses are based, and continuous consulting.

Many of INPUT's professional staff members have nearly 20 years' experience in their areas of specialization. Most have held senior management positions in operations, marketing, or planning. This expertise enables INPUT to supply practical solutions to complex business problems.

Formed in 1974, INPUT has become a leading international planning services firm. Clients include over 100 of the world's largest and most technically advanced companies.

Offices

NORTH AMERICA

Headquarters
1943 Landings Drive
Mountain View, CA
94043
(415) 960-3990
Telex 171407

Detroit
220 East Huron
Suite 209
Ann Arbor, MI 48104
(313) 971-0667

New York
Park 80 Plaza West-1
Saddle Brook, NJ 07662
(201) 368-9471
Telex 134630

Washington, D.C.
11820 Parklawn Drive
Suite 201
Rockville, MD 20852
(301) 231-7350

EUROPE

United Kingdom
INPUT, Ltd.
Airwork House
35 Piccadilly
London, W1V 9PB
England
01-439-8985
Telex 23116

France
La Nacelle
Procedure d'abonnement 1-74
2, rue Campagne Premiere
75014 Paris
France
322.56.46
Telex 220064 X5533

Italy
PGP Sistema SRL
20127 Milano
Via Soperga 36
Italy
Milan 284-2850
Telex 310352

Sweden
Athena Konsult AB
Box 22232
S-104 22 Stockholm
Sweden
08-542025
Telex 17041

ASIA/AUSTRALIA

Japan
ODS Corporation
Shugetsu Building
No. 12-7 Kita Aoyama
3-Chome Minato-ku
Tokyo, 107
Japan
(03) 400-7090
Telex 26487

K.K. Ashisuto
Daini-Suzumaru Bldg., 6th Floor
8-1, Nishi Shimbashi
3-Chome Minato-ku
Tokyo, 105, Japan
(03) 437-0654
Telex 781 26196

Singapore
Cyberware Consultants (PTE) Ltd.
2902 Pangkor
Ardmore Park
Singapore 1025
734-8142

INPUT
Planning Services For Management

INPUT® 1943 Landings Drive, Mountain View, CA 94043, (415) 960-3990