1989

An Analysis of the theoretical foundations of expert system technology

Leslie Anne Phillips
San Jose State University

Follow this and additional works at: https://scholarworks.sjsu.edu/etd_theses

Recommended Citation

Phillips, Leslie Anne, "An Analysis of the theoretical foundations of expert system technology" (1989). *Master's Theses*. 3102.

DOI: <https://doi.org/10.31979/etd.uk4y-wmgw>
https://scholarworks.sjsu.edu/etd_theses/3102

This Thesis is brought to you for free and open access by the Master's Theses and Graduate Research at SJSU ScholarWorks. It has been accepted for inclusion in Master's Theses by an authorized administrator of SJSU ScholarWorks. For more information, please contact scholarworks@sjsu.edu.

INFORMATION TO USERS

The most advanced technology has been used to photograph and reproduce this manuscript from the microfilm master. UMI films the text directly from the original or copy submitted. Thus, some thesis and dissertation copies are in typewriter face, while others may be from any type of computer printer.

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleedthrough, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send UMI a complete manuscript and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.

Oversize materials (e.g., maps, drawings, charts) are reproduced by sectioning the original, beginning at the upper left-hand corner and continuing from left to right in equal sections with small overlaps. Each original is also photographed in one exposure and is included in reduced form at the back of the book. These are also available as one exposure on a standard 35mm slide or as a 17" x 23" black and white photographic print for an additional charge.

Photographs included in the original manuscript have been reproduced xerographically in this copy. Higher quality 6" x 9" black and white photographic prints are available for any photographs or illustrations appearing in this copy for an additional charge. Contact UMI directly to order.

University Microfilms International
A Bell & Howell Information Company
300 North Zeeb Road, Ann Arbor, MI 48106-1346 USA
313/761-4700 800/521-0600

Order Number 1337336

**An analysis of the theoretical foundations of expert system
technology**

Phillips, Leslie Anne, M.S.

San Jose State University, 1989

U·M·I

300 N. Zeeb Rd.
Ann Arbor, MI 48106

AN ANALYSIS OF THE
THEORETICAL FOUNDATIONS OF
EXPERT SYSTEM TECHNOLOGY

A Thesis

Presented to

The Faculty of the Division of Technology
San Jose State University

In Partial Fulfillment
of the Requirements for the Degree
Master of Science in Special Major:
Industrial Technology

Leslie Anne Phillips

May, 1989

APPROVED FOR THE DIVISION OF TECHNOLOGY

Dr. Seth Bates

Dr. Ralph Bohn

Donald Betando

APPROVED FOR THE UNIVERSITY

Abstract

This is a philosophical study of the theoretical foundations of expert systems. The purpose of this study is to assess the relationship between cognitive models and expert system technology. In other words, how do the cognitive models assumed by researchers in expert system technology influence their work and conclusions? How closely do their projects in artificial intelligence actually replicate human intelligence?

The conclusions of this study are as follows:

- 1) The design of any expert system necessarily involves the incorporation of implicit or explicit epistemological concepts.
- 2) At this time there appear not to be any clear-cut, unquestioned answers as to precisely how the mind works or exactly what elements constitute human intelligence.
- 3) Some psychologists and philosophers contend that expert systems are necessarily missing many traits integral to human intelligence.

Table of Contents

	Page
Abstract.	iii
Chapter	
1. Introduction.	1
Statement of Purpose.	1
Statement and Significance of the Problem.	2
Limitations of the Research.	3
2. The Capabilities of Expert Systems: One View.	4
Expert System that Replicates the Subconscious.	4
The Computer as Guru	5
Expert Systems as Friends	7
Expert Systems Clone Mind	9
Intelligent Machine Plays Tic-Tac-Toe	13
"Renaissance Man" of Expert Systems	13
Expert Systems Replacing Experts	14
Computers with Emotions	15
Is Everything Automatable?	16
Summary	17
3. A More Critical Perspective	18
4. Analysis and Conclusions	32
Reference List	37

Chapter I

Introduction

Statement of Purpose

The researcher's interest in the subject of the philosophical suppositions of expert system technology was piqued while taking a graduate course at San Jose State University called "Knowledge Engineering Systems". During the course of the semester, a startling article entitled "Computers with Emotions" by L. Stevens (1987) was discussed. In the article, Stevens contends that computers could display emotions.

In order to determine how advanced expert system technology actually was, the researcher decided to consult experts from artificial intelligence centers at Hewlett Packard, Lockheed, FMC, Underwriters Laboratories, Xerox, and Stanford Research International. Through this research it was found that expert systems were truly remarkable from a technical perspective; they enable individuals to access phenomenal amounts of information and perform statistical, cost, and qualitative analyses regarding manufacturing operations, materials testing, electrical product testing, military system design, and cost calculation in a relatively painless manner. Such tasks would be substantially more cumbersome and time-consuming without the use of expert systems.

However, it did not seem clear whether replacing experts with expert systems would, in all cases, be wise. Therefore, the notion that expert systems are superior to experts in all respects was questioned. It was at that point the researcher decided to dedicate this thesis to the exploration of the theoretical foundations of expert system technology.

Statement and Significance of the Problem

Expert system technology is rapidly expanding, and is being implemented in American corporations almost as fast. With an ever-increasing number of areas of specialization and the added pressure of international competition, experts often seem to be in short supply.

But are we putting this technology to work all too quickly, without critically examining the proper role and limitations of expert systems? Can machines actually think? Is the 'expertise' of an expert system actually comparable to that of a human expert; that is, can computer software capture and duplicate the essence of human expertise? Is artificial intelligence in any way inferior to human intelligence? Do any aspects of human intelligence defy programming?

The International Joint Conference on Artificial Intelligence foresees a day when intelligent machines may not only provide instruction or intellectual stimulation,

but even social conversation and companionship. Some computer scientists expect that eventually computers will replicate all aspects of human intelligence. Is this possible?

Very enthusiastic, and not uncommon views, about the capabilities of expert systems will be explored in Chapter II. A more critical view, which details possible differences between human and machine intelligence, is provided in Chapter III. Although some claims made in Chapters II and III are quite arguable, most criticism and analysis will be reserved for Chapter IV. Chapter IV provides an analysis of the contents of Chapters II and III, in light of the philosophical and psychological assumptions implicit in each perspective.

Limitations of the Research

This study is limited to the analysis of artificial intelligence as embodied in expert systems. No investigation of other branches of artificial intelligence, such as robotics or machine vision, is undertaken.

Though the literature search for this study was not restricted geographically, all experts interviewed were working in the Santa Clara Valley of California.

Chapter II

The Capabilities of Expert Systems: One View

An Expert System that Replicates the Subconscious

Who would have ever thought computers would be capable of reproducing the processes of the subconscious mind?

Stanley J. Reiners is Vice President of Research and Development for Syndetic Corporation, a small software firm based in Omaha, Nebraska. Reiners believes that not only is it possible to make computers see, talk, and listen like human beings, but that they can be made to think like people, right down to the subconscious processes of the human mind (Myers, 1986, p. 34). He has been working on this project since 1972. Reiners, trained in nuclear physics at South Dakota State University, claims to have designed algorithms that can make a mainframe computer reproduce the processes of the subconscious.

He first came up with the idea of reproducing the subconscious while working as a subcontractor to NASA in Hampton, Virginia. He was working on the concept of interchangeable languages, "computer languages that would speak to each other, that would make it possible to switch from one language to another within a program, to make it possible for a programmer to use the best language for his needs at all times" (Myers, 1986, p. 34). In the process Reiner developed a learning model. The model was a static, manual model of thinking processes. A year later he

applied the same principles to a more accurate, dynamic, automatic, three-dimensional model that he eventually incorporated into a computer program for NASA. Reiner recounts, "We didn't realize what we were doing, but we programmed in learning. We identified that true knowledge had to incorporate learning because knowledge is not static but a state of equilibrium, a state that constantly changes" (Myers, 1986, p. 37). This discovery prompted Reiner to study psychology; he wanted to gain a better understanding of the ramifications of his program. He claims that the writings of Carl Jung gave him the most clear insight into the role of the subconscious in learning. Speaking of his own learning model concept, Reiner says, "We're here to declare the baby is born. I feel there is no area of computer usage in which this couldn't offer significant improvement" (Myers, 1986, p. 37). Reiner is confident that applications programs that learn as they execute, hold unimaginable potential for sophisticated expert systems.

The Computer as Guru

Another expert system about which remarkable claims are made is "Guru." Devised by Micro Data Base Systems Incorporated (MDBS), this natural-language-based artificial intelligence software allows users to use expert knowledge and reasoning for solving complex problems. Guru

incorporates standard business tools such as database management, spreadsheet analysis, text processing, business graphics, and remote communications. According to Gary Koehler, President and CEO of MDBS, "The program allows expert system applications to use information in the database or results of a spreadsheet calculation to draw conclusions" (Garretson, 1985, p. 4). The conclusions are made based on sets of rules, supplied as if-then statements, by the user.

Rule sets can be developed for all sorts of applications, from evaluating employment applications to determining least-cost routing in shipping. Users then provide information on specific cases, which allows the expert system to make decisions. For example, given certain criteria and the related data, the employment-evaluation system could decide whether or not to hire any number of candidates. Applications of any complexity (incorporating any number of rules) can be handled by the expert system, providing the product is run on adequate hardware. MDBS suggests using a PC with at least 512K bytes of memory and a minimum of 5M bytes of hard disk storage.

Guru is capable of using inductive and deductive reasoning. It can determine the circumstances required to achieve a desired result, or determine the result of a

given set of circumstances. This expert system works with compound and complex rules, as well as uncertain information. A percentage certainty can be assigned to information contained in a given rule. Based on the percentage certainty of the information, the expert system will include in its output a percentage of certainty for the conclusion.

Four different user interfaces are supported by the system to provide for various levels of user ability. Specifically, it allows for natural-language, menu-driven, command-driven, and programming-language interfaces. Guru integrates productivity tools such as a relational database-management system, a spreadsheet system, and a data inquiry system patterned after IBM's Structured Query Language. Spreadsheet cells can be accessed directly or can be defined as a result of an expert system consultation. In addition, rule sets can gain information from other rule sets. Guru can be run on local area networks such as Novell's Netware, 3Com's Ethershare, and IBM's PC Network.

Expert Systems as Friends

Daniel Hillis, an outspoken computer programmer, also has a very positive outlook on the power and potential of expert systems. Inspired by Marvin Minsky (the former

director of MIT's pioneering Artificial Intelligence Laboratory), Hillis, along with some sixty colleagues, has set out to build a thinking machine. Working for a company called, not-so-modestly, Thinking Machines Corporation, Hillis believes parallel processing is the key to artificial intelligence. He believes prior AI research did not succeed because of the computers that were used (Rothberg, 1985, p. 216). The more they learned, the "stupider" the computers got (in the sense that they got slower). Hillis solved this problem by breaking what computer scientists refer to as the Von Neumann bottleneck --the separation of the processor and the memory inside the computer, which allows the computer to process only one element at a time. Needless to say, processing only one element at a time slows the computing process. Hillis combined the processing and memory functions into a single, much faster, structure called parallel processing. This multiprocessing approach is likened to the neurocircuitry of the brain.

Hillis says about this revolution in thinking about the mind: "I think it will bother people, the idea that other things besides humans can be intelligent. But I think when it's all over, it won't present a threat to what's good about us" (Rothberg, 1985, p.216). Unfortunately, he does not explain this statement, but goes

on to recount his experience working with a "thinking, feeling" machine at M.I.T.. He viewed the computer as his friend. He cried when the big AI machine was decommissioned.

Expert Systems Clone Mind

"Expert systems clone one person's mind for ten, 100, or 1,000 people to use" (Guyer1, 1986, p. 30). These "cloned minds" or expert systems have been implemented at successful companies including General Electric, Proctor & Gamble, GTE Corporation, General Motors, Campbell Soup, and International Business Machines Corporation. GTE's expert system, called COMPASS, tracks 320,000 telephone circuits and notes missed connections all day long. COMPASS is a much quicker decision-maker than the engineers who previously diagnosed errors in the telephone service. This technology took less than a year to make its way from the research and development labs at GTE to complete implementation.

In addition to COMPASS, about 100 custom-built intelligent systems and in excess of 1,000 off-the-shelf systems are now in place at a variety of corporations. The corporate giant IBM, which had long been noncommittal about expert system technology, now admits to working on seventy expert systems for internal use and also plans on becoming a major supplier. Herbert Schorr, in charge of IBM's

artificial intelligence operations says, "Expert system technology is here now. We find applications wherever we look" (Guterl, 1986, p. 30).

The implementation of expert systems frees managers and supervisors from routine decision-making, resulting in cost and labor savings. Increased productivity and improved customer service are also made possible. Edward Mahler, manager of the AI program at E.I. du Pont de Nemours & Co., insists, "It's an extremely useful technology because it allows us to solve a class of problems that we couldn't solve before. It's saving us a hell of a lot of money" (Guterl, 1986, p. 31).

Digital Equipment Corporation reports savings in labor costs at \$25 million a year made possible through the use of an expert system that converts product orders into precise engineering parts specifications.

Westinghouse Corporation is one of the leading users of expert systems. Productivity at Westinghouse is way up due to the use of a system that automates the engineering of mechanical components. This system allows engineers merely to sketch the product onto a special electronic tablet, rather than producing stacks of drawings from scratch for each of the numerous products. The expert system then takes over, drawing on a wealth of information about the parts, to output engineering drawings and

specifications. Westinghouse Director of Research and Development, Andy Szabo, says, "Using this system is like printing money" (Guterl, 1986, p. 31).

Such engineering wizardry is viewed as a competitive weapon, allowing manufacturers the development time for new products. General Electric's Delco Product Division is taking full advantage of expert system technology, totally automating its design of electric motors. Stephen Dourson, Engineering Manager, claims using the system to design the parts automatically allows engineers to perform a task that would have normally taken four weeks in less than one hour. He additionally says, "It's a competitive world out there, and we want to be able to come back to a customer after two days and hand him four or five different designs to choose from" (Guterl, 1986, p. 32).

Diagnostic expert systems are also very valuable to industry. For example, Campbell Soup uses one such system to diagnose problems in a cooker that sterilizes soup after canning so a specialist doesn't have to be called. Ford Motor Company uses a diagnostic expert system in manufacturing to quickly identify malfunctioning robot arms. The use of smaller diagnostic expert systems to troubleshoot equipment in the field can reduce service calls. All General Motors dealerships have access to an expert system that identifies engine problems without

having to use a service agent. Motorola employs an expert system to diagnose problems with customers' Unix computers. IBM is working on distributing a system that will allow field service engineers to test disk drives more easily.

Companies have two basic options with regard to expert systems. They can either use completely custom-made systems or they can bring in expert system shells. The shells are cheaper than starting from scratch and they can get them up and running faster. Hence, many high-tech firms are buying expert system skeletons, or shells, these days. They can build a relatively simple, but efficient, expert system for under \$100K, including the cost of training a programmer.

Du Pont is a major corporate user of expert system shells. They have approximately 150 such systems in use, performing various tasks, from selecting the best kind of rubber for customers, to scheduling machines on the factory floor and diagnosing equipment malfunctions. Regarding their decision to use shells rather than starting from scratch and learning to make their own systems, Ed Mahler of du Pont says, "We refused to put the company through a culture change and learn all the artificial intelligence wizardry" (Guterl, 1986, p. 36). Du Pont has no regrets about choosing to use expert system shells.

The du Pont training program claims to train twenty experts a week, each capable of fully testing systems within a month. The cost is approximately \$20,000 per person, including equipment. Mahler says, "The cost for du Pont has been peanuts. The last ten systems we started have given us a 800% return on investment" (Guterl, 1986, p. 36).

An Intelligent Machine that Plays Tic-Tac-Toe

An expert system that plays tic-tac-toe is on display at the Boston Museum of Science. The computer program relies on the judgment of seven internal "experts" to play the game (Nash, 1987, p. 78). For example, WIN searches for winning moves. DEFEND blocks wins by the opponent. DECIDE weighs the suggestions of the other six experts. The machine has a high success rate.

The "Renaissance Man" of Expert Systems

Implementing expert systems in factories can be difficult because very few manufacturing operations involve only one realm of expertise (Smith, 1987, p. 141). It's hard to get two experts in one field to agree, let alone a group of experts from different disciplines. But Major Steven R. LeClair, head of research in artificial intelligence for manufacturing at the Materials Laboratory at Wright-Patterson Air Force Base, decided it was necessary

to undertake the task. The result is what LeClair calls a Multiexpert Knowledge System (MKS), or what could be viewed as the "renaissance man" of expert systems.

MKS has made remarkable accomplishments. For example, it discovered the most efficient requirements for curing complex plastic composites. Previous requirements used by the aerospace industry involved baking a 256-layer, graphite lamination to be used for airframe parts for twelve hours, whereas MKS developed an equally effective, more timely scheme for curing the composite in only three hours. Some experts were sceptical at first, but the process really works. LeClair hopes to achieve similar results in other process-control applications in the future.

Expert Systems Replacing Experts

Joseph Kroger, Vice-Chairman of Unisys Corporation, expects very widespread use of artificial intelligence in expert systems by the year 2000 (Kroger, 1987, p. 38). He believes that computers can be made to manipulate symbols just the way human beings do, only much faster. For example, NASA scientists spent eight years trying to develop a way to eradicate carbon dioxide from space shuttles; this task was accomplished in just four weeks by rapid computer prototyping.

Another application of this technology involved optimizing seat revenue; an expert system was used to analyze various factors related to seating for a major U.S. airline. As a result, the airline gained increased profit, more efficient scheduling of personnel and facilities, and a competitive advantage. Expert systems can also be used to schedule manufacturing for companies or to diagnose printed-circuit-board failures. Potential applications are seemingly endless. The relevant knowledge of virtually any expert can be readily stored, accessed, and used by computers.

Computers with Emotions

Lawrence Stevens promotes the idea that emotions are based on rules (Lawrence, 1987, p. 39). Hence, emotions can be programmed into the computer as if-then statements. Stevens claims that even sudden personality changes, such as those brought on by prayer or meditation, are rule-based.

Imagination may also be able to be programmed into computers. The ability of the computer to make guesses about what would happen under a certain set of circumstances can be viewed as imagination. According to Stevens, computer programs could even use pleasure or satisfaction as a way of deciding what to do.

The computer could exhibit pleasure by exploring imaginative possibilities in detail and by working out possible variations and ramifications. If a ramification produces a negative emotion, the program would--as a human would--explore another possibility that might have more "satisfying" consequences. A program with "imagination" could be said to be a very efficient daydreamer, imagining for the sake of pleasurable solutions (Stevens, 1987, p. 39).

In short, Stevens believes computers can be given imaginations, emotions, and feelings comparable to those of human beings.

Is Everything Automatable?

Apparently some managers are resisting the implementation of expert systems. Expert systems could be implemented in management applications and processes, but they are often not ("Expert systems ready...", 1987, p. 41). This is due to the attitude that "everything is automatable up to but not including what I do" (Expert systems ready...", 1987, p. 41). In other words, the lack of use of expert systems in management applications is being blamed entirely on the arrogance of those in charge, rather than admitting possible shortcomings or limitations of expert systems.

Summary

Many people believe those working in expert systems technology have enjoyed uncanny success in replicating various aspects of human intelligence. Researchers claim to have created intelligent machines that have accomplished everything from controlling complex manufacturing processes to displaying human emotions and wisdom.

Chapter III

A More Critical Perspective

Some AI researchers are convinced computers may never really think like people. The Dreyfus brothers have a thought-provoking, detailed explanation of this matter (Dreyfus & Dreyfus, 1986, pp. 42-61). The Dreyfuses are both professors. Hubert Dreyfus is a professor of philosophy, previously at Massachusetts Institute of Technology. Stuart Dreyfus is a professor of industrial engineering and operations research at the University of California, Berkeley. They say that leading computer scientists have long been hoping to develop thinking machines that are not dependent upon human control. Present AI researchers believe they are very close to doing so. One of the most prominent AI researchers from MIT, Marvin Minsky, says, "Today our robots are like toys. They do only the simple things they are programmed to. But clearly they're about to cross the edgeless line past which they'll do the things we [humans] are programmed to [do]" (Dreyfus & Dreyfus, 1986, p. 42).

Encouraged by such positive statements, the U.S. Department of Defense is investing millions of dollars into perfecting this technology, hoping that they will eventually gain completely autonomous war machines that will respond to threats without human intervention. Top business executives are sinking substantial amounts of

money into expert systems, hoping that some day the wisdom of the expert systems will equal or top that of their key managers. AI entrepreneurs predict that intelligent systems will soon outperform human beings in the classroom, in the home, and at work.

However, the Dreyfuses believe that computers will never live up to all these expectations. "After 25 years of research, AI has failed to live up to its promise, and there is no evidence that it ever will. In fact, machine intelligence will probably never replace human intelligence simply because we ourselves are not 'thinking machines.' Human beings have an intuitive intelligence that 'reasoning' machines simply cannot match" (Dreyfus & Dreyfus, 1986, p. 44).

The Dreyfuses hope civilian managers and military officials see the shortcomings and refrain from using such machines, but contend that they may be blinded to the reality of the situation by their enormous hopes and the fear of having wasted the large sums of money invested thus far. Computers cannot come close to matching the expertise of seasoned business managers, master teachers, or skilled air traffic controllers. Because of this, these individuals should not be replaced by computers. In addition, "Computers that teach and systems that render 'expert' business decisions could eventually produce a

generation of students and managers who have no faith in their own intuition and expertise" (Dreyfus & Dreyfus, 1986, p. 44). Computers are quite capable of the most sophisticated calculations, but are of their very nature seriously lacking in judgment.

Acquiring human know-how is no simple matter for computers. How can know-how be stated for computers to understand? For example, though most people know how to ride a bicycle, they would be hard pressed to formulate precise rules to teach anyone else how to do it. There is a very fine line between the feeling of falling over and the sense of being a little off balance when turning. And how we would respond to a certain wobbling feeling is quite uncertain until the situation occurs. Yet know-how, acquired from practice and often painful experience, enables us to ride a bicycle.

Countless other aspects of daily living are equally difficult to reduce to "knowing that." "Know-how" is necessary to carry on an appropriate conversation with strangers, friends, and family in numerous contexts, including the street, a party, or the office. Walking is a simple function for most people, yet the mechanics of walking on two feet is so complicated that engineers cannot begin to reproduce them in artificial devices.

Such know-how is not innate, as is a bird's skill at building nests. It is a learned function. It is learned through a process of trial and error, as watching any small child begin to walk will demonstrate. Imitating those walking around them can also help toddlers gain the skill. Adults also require instruction and experience to acquire skills. "Knowing that" is knowledge guided by rules. Know-how is experience-based.

People usually pass through five skill levels in attaining know-how: novice, advanced beginner, competent, proficient, and expert.

During the first of these levels, the novice stage, people generally learn the facts related to a particular skill and the rules for action that are derived from those facts. For example, people learning to operate a stick shift automobile are told at what speed to change gears and at what distance (depending on the speed) to follow other cars. These rules do not consider the context, such as the number of stops the driver needs to make or the density of the traffic.

Likewise, novice chess players learn a formula which allows them to assign pieces point values regardless of their position. They learn the rule: "Always exchange your pieces for the opponent's if the total value of the pieces captured exceeds that of pieces lost" (Dreyfus & Dreyfus,

1986, p. 45). Novices generally do not realize that this rule should be violated under certain circumstances.

Novices progress to the advanced beginner stage after a significant amount of experience. Advanced beginner drivers are aware of situational elements that are not yet objectively defined. For instance, they are more aware of engine sounds when switching gears. Advanced beginner chess players know to avoid overextended positions. They can also pick up on situational clues like a strong pawn structure or a weakened king's side. In all such circumstances, experience is much more important than any sort of verbal description.

The training wheels on a child's first bicycle are like initial rules which allow beginners to accumulate experience. But rules must be put into perspective in order to proceed. For example, at the competent stage, drivers do not merely follow rules; they drive with an aim in mind. If they want to get from point A to point B quite quickly, they choose their path with the traffic in mind, paying little attention to passenger comfort. They enter traffic more daringly, follow other cars more closely than they "should," and may even break the law. A competent chess player may ignore the lessons they learned as a beginner and accept some personal losses in order to strategically attack their opponent's king.

A major difference between beginners and more competent individuals is their level of involvement. Novices and advanced beginners do not usually feel entirely responsible for what they are doing because they are only applying learned rules; if they make a mistake, they can blame the rules instead of themselves. On the other hand, competent performers, who have a goal and a plan for achieving it, often feel greater responsibility for the outcome of their choices. A successful outcome can bring great satisfaction, while disasters may not be easily forgotten.

Those learning a new skill make conscious decisions after considering various options. But this detached, deliberate, painstaking model of decision making is the exception rather than the rule in our everyday lives. Proficient performers do not need to rely on deliberation. Rather, memories of similar past experiences prompt plans like those that worked before. Proficient performers are able to recall and apply whole situations without having to break them down into rules or components. For instance, a boxer knows the moment to begin an attack not by mentally reciting and following rules. Rather, the whole scene triggers the memory of a successful past attack. The boxer is using his know-how, or intuition.

Intuition is not an unconscious, child-like recognition of new situations similar to remembered ones. In other words, we advance from analytic behavior to skilled behavior. Conversely, young children only understand concrete examples at first, then gradually learn abstract concepts. Adult intelligence may be misunderstood so often precisely because this pattern in children is so well known.

There is more to intelligence than mere calculative rationality. As a matter of fact, experts who attempt to reason things out may regress to the level of novice, or perhaps, competent performer. This does not mean, however, that deliberate rationality plays no role in intelligence. Sometimes detached deliberation may keep a person from falling victim to tunnel vision. In other words, focusing on certain seemingly trivial details of a situation may allow another perspective to come to mind.

The story of an Israeli fighter pilot demonstrates this point. Just after vanquishing an expert opponent, he found himself confronted by another fighter from the enemy squadron who appeared to be brilliantly executing a series of masterful ploys. Things were looking quite grim for the Israeli pilot until he abandoned his intuition and deliberated. This allowed him to realize that the surprising maneuvers of his opponent were no more than the

rule-following, predictable behavior of a beginner. This insight allowed him to defeat the pilot.

But intelligence is not necessarily dependent upon facts. Digital computers, which are essentially complex structures of on-off switches, were first used to make scientific calculations. During the late fifties, however, two computer researchers, Allen Newell and Herbert Simon, began to toy with the idea that general symbols could be manipulated by computers. They realized that symbols could be used to express elementary facts about the world and rules to represent the relationships between the facts. Initially, Simon and Newell theorized that computers programmed with such rules and facts could recognize patterns, comprehend stories, solve problems, and do anything else an intelligent human being could do. But they quickly realized that crucial aspects of problem solving were missing. For instance, the computers could not separate relevant and irrelevant operations. Accordingly, the programs could do little more than solve problems and prove theorems of logic.

Newell and Simon's approach was abandoned by the late sixties. At that time, researchers at M.I.T. began to focus on processing methods rather than attempting to copy reports of methods people claimed they used to solve problems. The researchers realized that the computer had

to somehow simulate understanding and intuition in order to solve real-world problems. Marvin Minsky describes the M.I.T. approach in the following manner:

If we ... ask ... about the common everyday structures --that which a person needs to have ordinary common sense--we will find first a collection of indispensable categories, each rather complex: geometrical and mechanical categories of things and of space; uses and properties of a few thousand objects; hundreds of "facts" about hundreds of people; thousands of facts about tens of people; hundreds of facts about hundreds of organizations ... I therefore feel that a machine will quite critically need to acquire on the order of a hundred thousand elements of knowledge in order to behave with reasonable sensibility in ordinary situations. A million, if properly organized, should be enough for very great intelligence (Dreyfus & Dreyfus, 1986, p. 49).

This approach, however, was no better than Newell and Simon's, in that each program worked only within its limited specialty and could not be applied to any other problems. The programs were also lacking in semantics. In other words, they didn't understand what their own symbols meant. For instance, the STUDENT program created by Daniel Bobrow, which was supposed to solve simple algebraic story

problems, interpreted the phrase "the number of times I went to the movies" as the product of the two variables "number of" and "I went to the movies." The program mistakenly thought "times" was a multiplicative operator linking those two phrases.

Joseph Weizenbaum, another computer science professor at M.I.T., wrote a program called ELIZA which vividly demonstrates how much apparent intelligence a computer can exhibit without having any real understanding. ELIZA is a program that imitates a therapist by using simple tricks such as turning statements into questions. When told "I'm feeling sad", it responds "Why are you feeling sad?" When the program has no stock response handy, it prints out statements like "Tell me about your father." Surprisingly, many people were easily fooled by these tricks. Weizenbaum was shocked to find people asking others to leave the room so they could divulge their deepest secrets to a computer.

On one occasion the shallowness of the computer was exposed unintentionally when someone typed "I'm feeling happy," and then went on to correct himself by typing, "No, elated." The computer responded, "Don't be so negative." In other words, it had been programmed to use that rebuke when there was a "no" in the input.

Within about five years, the shallowness of programs like Minsky's became apparent. Computer scientists were

forced to take a new approach. In 1970, Minsky and another professor, Seymour Papert, decided to deal with isolated sub-worlds and gradually build upon them, since trying to tackle common-sense knowledge all at once was too great a task. Shortly thereafter, Terry Winograd made a computer program that advanced artificial intelligence significantly by getting computers to understand natural language (human language, rather than machine language). The program was called SHRDLU. A robot arm that could move a group of variously shaped blocks was simulated on a TV screen. SHRDLU allowed people to engage in a conversation with the computer, making statements, giving commands, and asking questions within the world of movable blocks. Facts about the blocks, semantics, and grammatical rules were all employed in this program. SHRDLU apparently understood language within its fixed domain.

Minsky and Papert attributed Winograd's success to his choosing a simple problem which easily lent itself to the restricted application or "microworld." The men believed that by integrating numerous microworlds they would be able to give real-life understanding to computers. The Dreyfuses have the following to say about the matter,

Unfortunately, this research confuses two domains, which we shall distinguish as "universe" and "world." A set of interrelated facts may constitute a

"universe" such as the physical universe, but it does not constitute a "world" such as the world of business or theatre. A "world" is an organized body of objects, purposes, skills, and practices that make sense only against a background of common human concerns. These "sub-worlds" are not isolable physical systems. Rather, they are specific elaborations of a whole, without which they could not exist. (Dreyfus & Dreyfus, 1986, p. 50)

If the subworlds created by these computer programmers were actually subworlds, they would not have needed to be broadened and linked to encompass the everyday world, since each one would have already incorporated it. Microworlds cannot be combined and extended to portray everyday life because they are isolated, meaningless domains. AI researchers worked with this model doomed to failure for about five years. Winograd himself later acknowledged, "The AI programs of the late sixties and early seventies are much too literal. They deal with meaning as if it were a structure to be built up of the bricks and mortar provided by the words" (Dreyfus & Dreyfus, 1986, p. 50).

AI has been struggling unsuccessfully with what is referred to as "the problem of common sense" since the late seventies. Finding a way to get computers to retain and

access all the facts people seem to know has alluded researchers all this time. Needless to say, Minsky and Simon's prediction made in the mid-sixties, that by the mid-eighties computers would be capable of doing everything human beings can do, never came to fruition.

Unfortunately, computers do not have a basic understanding of human life. As a result, it is very difficult, if not impossible, for them to interact intelligently with people. People understand all sorts of things simply by virtue of being human. For instance, we realize that moving physically forward is generally easier than moving backward, that insults make us angry, and the list goes on. Programming all of this into a computer as rules and facts is unimaginable. AI workers describe this task as giving computers our belief system. One faulty assumption of this project is that our beliefs are easily gathered and stored as facts.

Even if a way to retain these facts could be found, computers cannot be programmed for context. For example, computers cannot be programmed to know that a car is simply going "too fast." The program would need to be free of interpretation. We must specifically include the fact that the car is traveling "25 miles per hour," for example. In addition, computers operate according to precise rules, such as "shift to second at 25 miles an hour," not common

sense rules such as "under normal conditions, shift to second at around 25 miles per hour."

Even if all facts could be compiled in a context-free form, the computer still wouldn't know what to do with them because it would be unable to draw only on the relevant facts or rules. There are exceptions to almost every rule, and exceptions to almost every exception. It would be nearly impossible to include all of them and somehow inform the computer to know which exception is to be used at any given time.

"In the final analysis, all intelligent behavior must hark back to our sense of what we are. We can never explicitly formulate this in clear-cut rules and facts; therefore, we cannot program computers to possess that kind of know-how" (Dreyfus & Dreyfus, 1986, p. 51). In short, computers will probably never be capable of thinking like people since they are, in so many ways, dissimilar from people.

Chapter IV

Analysis and Conclusions

How can the differing perspectives on the intelligence of expert systems be accounted for? Do expert systems display human-like intelligence? Is there a clear-cut answer?

One important realization is that all theories are not written from the same meaningful perspective (Rychlack, 1981, pp. 20-24). As is pointed out by James Brule, investigating the philosophical and psychological foundations of expert systems can provide much insight into the actual potential and accomplishments of this technology (Brule, 1986, pp. 16-18).

What is thinking? What does it mean to "know" something? What is meaning? What is intelligence? How can these topics be intelligently discussed as they relate to expert systems if the matters themselves are not first critically examined and understood? These questions are central to psychology and philosophy (especially epistemology), but have just as much bearing on expert system technology (Waldrop, 1987, pp. 12-20). For example, someone developing an expert system must decide how knowledge will be represented in the system. Some notion of what intelligence is, whether implicit or explicit, necessarily precedes any attempt to program intelligence into a system. Though there are no absolute, undebatable

answers to most of the critical questions, the assumptions of the theories adopted should be fully understood (made explicit) since they have important consequences.

Traditional psychology, which is generally the psychology borrowed by those working in artificial intelligence, bases its approach on a method used in natural science (Rychlack, 1979, p. 38). Thus, human intellect and behavior are explained in terms of cause and effect. "Mental processes are reduced to empirical phenomena observable within a causal chain" (Stewart, 1974, p. 119). The mind is, basically, viewed as a machine.

Humanistic psychology, on the other hand, refutes the notion that all mental processes are caused by physical events. Humanistic psychology charges that "mechanistic psychology omits the most important aspect of human behavior, namely, its meaning" (Stewart, 1974, p. 119).

S. Reiners claims his expert system which "reproduces the processes of the subconscious" is in line with Jungian psychology. But is it really? Both men believe that knowledge is not static; learning is critical. But the analogy ends there. What is learning? What is knowledge? Reiners' program "learns" strictly by being programmed with new, clearly defined facts and formulas. Jung's psychology, on the other hand, is based on human experience. Unlike computers, human beings are capable of

experiencing pain, compassion, fear, love, and, joy.

Moreover, in Jung's view, the subconscious is much more than a complex machine. Jung dedicated his life to exploring the spiritual, inner world. As spiritual beings, we have the capacity for self-transcendence, for the realization of purpose and meaning.

Much of what we know as human beings defies categorization and precise explanation. Many researchers in expert system technology might do well to learn a lesson from the following five humanistic postulates:

1. Man, as man, supersedes the sum of his parts (that is, man cannot be understood from a scientific study of part-functions.)
2. Man has his being in a human context (that is, man cannot be understood by part-functions which ignore interpersonal experience.)
3. Man is aware (and cannot be understood by a psychology which fails to recognize man's continuous, many-layered self-awareness.)
4. Man has choice (man is not a bystander to his existence; he creates his own experience.)
5. Man is intentional (man points to the future; he has purpose, values, and meaning.) (Yalom, 1980, pp. 18-19)

Computers are quite capable of replicating certain aspects of human intelligence, especially those involving logic and calculation, but are gravely deficient in other areas. Computers have no common sense, no emotion, no psyche, no spiritual nature. They have no free will, no experience. No number of rules and specifications can breath life and understanding into an expert system. The majority of the human traits most valued by humanistic psychologists cannot be programmed into a computer.

Humanistic Psychology is primarily concerned with those human capacities and potentialities that have little or no systematic place ... in behaviorist theory: e.g., love, creativity, self, growth, organism, basic need-gratification, self-actualization, higher values, being, becoming, spontaneity, play, humor, affection, naturalness, warmth, ego-transcendence, objectivity, autonomy, responsibility, meaning ... transcendental experience, psychological health, and related concepts.

(Yalom, 1980, p. 18)

Computers are great at logic and following rules, but most of the traits most valued by humanistic psychology cannot be programmed into a computer. A humanist would hold that a computer necessarily misses out on much of what it means to be an intelligent human being. So who's right and who's wrong? What is the status of expert system technology? It

is clear that computers are very efficient in controlling industrial processes and performing mechanistic functions. Computers certainly allow people to access large amounts of information. But are expert systems truly worthy of the name "expert"? Isn't there a more appropriate term? Does not human expertise entail much more? I believe it does.

References

- Brule, J.F. (1986). Artificial intelligence theory, logic, and application. Blue Ridge Summit, PA: Tab Books, Inc..
- Dreyfus, H. & Dreyfus, S. (1986, January). Why computers may never think like people. Technology Review, 89(1), 42-61.
- Expert systems ready, but are managers? (1987, June 8). Design News, 43(11), 41-42.
- Garretson, R. (1985, November 5). Expert advice contained in 'guru' package. PC Week, 4.
- Guterl, F. (1986, October). Computers think for business. Dun's Business Month, 128(4), 30-37.
- Kroger, J. (1987, July/August). Artificial intelligence: a new reality. Futurist, 21(4), 38-40.
- Myers, E. (1986, December 15). To Jung, with thanks. Datamation, 32(24), 34-37.
- Nash, J. (1987, February 9). Expert systems. Time, 129, 78.
- Rothenberg, R. (1985, December). Daniel Hillis on artificial intelligence. Esquire, 104(6), 214-216.
- Rychlack, J.F. (1979). Discovering free will and personal responsibility. New York: Oxford University Press.
- Rychlack, J.F. (1981). Personality and psychotherapy. Boston: Houghton Mifflin Company.
- Smith, E. (1987, May 11). The 'renaissance man' of expert systems? Business Week, (2998), 141.

- Stevens, L. (1987, July/August). Computers with emotions.
Futurist, 21(4), 39.
- Stewart, D. (1974). Exploring phenomenology. Chicago:
American Library Association.
- Waldrop, M.M. (1987). Man-made minds. New York: Walker and
Company.
- Yalom, Y.D. (1980). Existential psychotherapy. New York:
Basic Books, Inc..