

AD-A111 895

RAND CORP SANTA MONICA CA
EXPERT SYSTEMS: NEWEST BRAINCHILD OF COMPUTER SCIENCE, (U)
JUN 81 A M SHOEN
RAND/P-6552

F/6 9/2


UNCLASSIFIED

NL

1 of 1
56
25-0000



END
DATE
FILMED
4-82
DTIC



MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

(2)

AD A111845

EXPERT SYSTEMS: NEWEST BRAINCHILD OF COMPUTER SCIENCE

Ann M. Shoben

June 1981

DTIC FILE COPY

DTIC
MAR 10 1982
A

P-6552

This document has been approved for public release

42 08 08 042

PREFACE

This paper is a journalistic account of a week-long workshop on expert systems held in San Diego in August 1980. Funds for the conference were provided by The National Science Foundation, the Advanced Research Projects Agency of the U.S. Department of Defense and The Rand Corporation. The author is deputy director of Rand's public information office.



Preparation for
MIA GRAAI
TAB
Assigned
Collection

Dist Special

A

The knowledge-base worker's a star
With followers local and far.
I'd be happier yet,
As would you, I will bet,
If only we knew what we are.

Peter Hart

EXPERT SYSTEMS: NEWEST BRAINCHILD OF COMPUTER SCIENCE

From a distance the blackboard sketch resembled a child's crude drawing of a frog or the outline of a squat and faintly comic space creature. It was instead a paradigm--the speaker's abstract representation of a "messy" real life problem. With labeled balloons and connecting lines, he was explaining to an audience of his peers the approach he and his partner were taking to a test problem involving oil and chemical spills on inland waters.

The group was small, about 50 men and women in whose heads reside most of what is known in this country about ^{Dr. Hart discusses} an emerging field of computer science called "expert systems." These are computer programs that seek to combine expert knowledge about a domain (in theory, any domain) with expert methods of conceptualizing and reasoning about that domain. The decision-making power of such programs rests upon a knowledge base that puts together factual information about the domain with the heuristics (informal "rules of thumb") experts use to rapidly find solutions to problems. This specialty in turn is part of

artificial intelligence--the science, as one definition has it, of making machines do things that people need intelligence of a high order to do.

Despite its small size and relative newness, the expert systems field is attracting widespread interest because of its demonstrated usefulness in medicine and chemistry and for its potential application to such diverse and complex problem areas as structural engineering, military strategy, air traffic control, crisis management and many others.

The first expert systems were created at Stanford University a little over a decade ago. The prototype, a program called DENDRAL, aids chemists in determining the molecular structure of unknown compounds. Because the molecular structure of a compound must be understood before its other properties can be studied--properties related to pharmacology or toxicology, for example--this program promises an important contribution to biomedicine. Other programs specialize in diagnosing and recommending therapy for patients with infectious diseases, blood disorders, say, or meningitis. When completed, such tools will be able to assist physicians working on difficult cases and paramedics in remote or medically underserved areas.

Held in August in San Diego, the conference that brought this group together for the first time was organized by Frederick Hayes-Roth and Donald Waterman of Rand's information sciences department, and by Douglas Lenat of Stanford University's computer sciences department. The National Science Foundation, the Advanced Research Projects Agency of the Department of Defense and Rand provided the funds.

Eight different approaches to the creation of expert systems were represented at the meeting. These approaches (usually themselves called "systems" or "tools") were developed at such centers of artificial intelligence research as Stanford, Carnegie Mellon University, the University of Southern California, Rutgers University, the University of Pittsburgh, Massachusetts Institute of Technology, SRI International, and Rand. The air was charged with eagerness and some anxiety because, despite the organizers' attempts to downplay it, the meeting had taken on something of the flavor of a contest.

The conference was organized along two converging lines. One dealt with the conceptual issues--definitions, framework, performance, evaluation--that bedevil any fledgling science, and the other with putting into practice the tools of the trade.

It was this latter exercise--the actual attempt during the course of the week to build not one expert system but eight--that was the focus of excitement and attention. Taking part in the experiment were eight two-man teams (one for each system) composed of a leader, the theoretician, and a "wizard," the name given to a computer programmer of extraordinary speed and competence. The tools these "knowledge engineers" brought with them were systems (with such names as EXPERT, HEARSAY, ROSIE, EMYCIN, OPS, PROSPECTOR, AGE, and RLL) that had been developed at their home institutions. Typically these tools have two parts: an analytical framework (paradigm) into which the problem must be fitted, and a language for programming. An exception was Rand's ROSIE, a natural programming language much like English, with English words, grammar and syntax, which allows the analyst to provide his own

paradigm. No one expected workable programs to emerge from the experiment--a task that typically requires months, even years, of effort. The attendees hoped instead that the fund of common knowledge would be increased and insights would be gained into the strengths and weaknesses of each of the eight systems.

Choosing the right problem for the exercise required careful thought. It had to be difficult enough to challenge the expert systems engineers, yet of not so vast a scale as to overwhelm them in the short time defined by the conference. Also, ideally, it should be generalizable enough to serve as a model for more complex problems of the same type.

The problem finally selected was a real one outlined by experts, Carroll K. Johnson and Sara R. Jordan of Oak Ridge National Laboratory in Oak Ridge, Tenn. Like all large facilities, Oak Ridge stores on its premises a variety of motor oils, lubricants and chemicals for housekeeping and maintenance purposes. Occasionally accidental spills occur. A typical spill might involve a pool of oil that leaked from a piece of construction equipment and washed into a storm drain during a rain shower.

Recently Oak Ridge undertook an aggressive emergency response program aimed at preventing, or failing in that, cleaning up the damage done to streams on its property by such accidents. The action was prompted in part by increasingly stringent government regulations (and stiff penalties) and in part by the belief that government agencies should not be the last to demonstrate concern for health and the environment. In this, however, Oak Ridge is only a few jumps ahead of a

federal law requiring all installations, whether public or private, large or small, with streams on their land to develop similar emergency procedures.

Johnson, a crystallographer with a hands-on, working knowledge of computers, was convinced that the spill problem--which he saw as similar in its overall outlines to a military command and control situation--would lend itself to the techniques and tools of an expert systems approach. With Jordan, who is a computer specialist, he prepared in advance some of the materials that would be needed for the exercise. These included a general description of the spill problem at Oak Ridge, the methods and procedures that had been devised for coping with it, several case studies of real and hypothetical spills, a partial inventory of the various oils and chemicals stored in the Oak Ridge complex, and diagrams of the surrounding drainage basin.

An accurate representation of the drainage system was crucial information for the problem solvers because, drainage basins being what they are, any fluid that is spilled in sufficient quantity is likely to find its way via ditches and storm drains into the surrounding waterways. And even a quart of oil can cause a visible sheen (an environmental no-no) on a body of water. At Oak Ridge, the spill route follows White Oak Creek, a small stream bounding one side of the property, into White Oak Lake, and from there into the Clinch River, which serves as a water supply for several communities.

Johnson's analogy of a military gaming scenario helped to put the various components of the exercise into perspective: A large defense network with a hierarchical command structure has been organized to

protect a vulnerable resource (the water supply). One defense company in the network (Oak Ridge) is assigned the responsibility of protecting one part of the resource (White Oak Creek). The villain (spill demon) can penetrate the outer defense (spill prevention program) with a sabotage column (spill flow). There are many infiltration trails (drainage routes), all leading toward the resource. The villain's objective is to disable (pollute) as much of the resource as possible. The defense network has a defense strategy (spill countermeasures plan) which must be implemented and tailored to the circumstances. This involves detecting the attack, notifying the chain of command, setting up a defense line, determining the strength and type of the villain's penetration, evaluating the situation, mounting a counterattack, cleaning up the resistance, handling the captured material and filing a report of the action.

Containing the spill once it has been spotted and tracing it back to its source are obviously only two aspects of a many-sided problem, but they are important and require immediate attention. The teams were instructed to take whatever approach to the exercise they wished-- to deal with the overall problem or to concentrate on only one or two aspects of it-- but they were each urged to focus on one common sub-problem so that a basis for comparing methods would be available at the completion of the experiment. The sub-problem was that of locating the source of the spill.

The goal of an expert systems approach is to encode in a computer program both the facts an expert has at his command and his methods of reasoning about them. So the first step in the process was an interview

session during which the knowledge engineers grilled Johnson and Jordan about details of the spill problem at Oak Ridge. In some respects this is the most arduous task facing both expert and examiner in the construction of an expert system--one requiring tact, patience and considerable mutual forbearance. Difficulties arise when the expert tries to make explicit what is often intuitive--as when a physician knows, without consciously going through the diagnostic procedure he learned in medical school, that a patient is suffering from mononucleosis and not Asian flu.

After two lengthy question-and-answer sessions with Johnson and Jordan, the teams went to their rooms to work. Here computer terminals had been installed with telephone links to the teams' home computers.

Although one member of each team was responsible for formulating the paradigm and the other for translating it into computer language, in practice the distinction was not so sharp. In each case the team members had worked closely together in the past, and like most people who know each other well, they argued amicably and at length in a kind of private shorthand, the role of one often seeming to shade into that of the other. All-night sessions, at which gallons of coffee and countless doughnuts were consumed, became the rule. Bleary-eyed wizards nodded through morning sessions, while their equally weary colleagues reported their progress to the group at large. Sometimes the first approach had to be abandoned and another tried. The teams compared notes, sought and gave advice, commiserated. From time to time, the experts were called in to clarify a point or go over a procedure again.

By mid-week, each of the teams had arrived at a pilot solution to the spill problem, or to those components of the problem on which they had chosen to concentrate. If these were not complete, fully operational programs, they nonetheless represented respectable first passes that the teams could live with.

By now, however, it had become clear that not all the systems were congenial to the type of problem represented at the conference. Those, like EMYCIN, EXPERT and PROSPECTOR, which are powerful tools for solving diagnostic problems (medical in the case of EMYCIN and EXPERT, and exploring for mineral deposits in the case of PROSPECTOR) were unsuited for the multi-level messiness of the oil spill exercise. However, by some tinkering and judicious force-fitting all of the systems were eventually adapted to handle one or more of the sub-problems.

The PROSPECTOR team, for example, selected the problem of giving advice to an inexperienced coordinator in the field about the best strategy to use in locating the source of the spill. They chose this task because PROSPECTOR is particularly good at weighing and balancing different considerations and can respond to volunteered data. Consequently, they quickly obtained a simple system that could respond to volunteered information, seek more information in an efficient way, weigh and balance competing factors in reaching conclusions and explain its conclusions.

Some systems had difficulty dealing with a dynamic situation where new information frequently requires the program to change its responses or re-order its priorities. For example, if the computer is in the midst of an elaborate backtracking procedure looking for the source of a chemical spill and word comes in from the field that someone has

stumbled on a leaking container of acid in Building A, the computer should be flexible enough to recognize that this is the likely cause of the trouble, interrupt the search procedure and go on to the next step in the process. But in systems designed to handle static problems, like diagnostic ones, where goals and priorities tend to be fixed, the usual reason for changing a program response is to correct erroneous information, not to process new information.

Those who chose to work on the overall problem were confronted by issues they had never faced before. One was how to coordinate a variety of tasks and the human agents who carry them out when the tasks often have competing priorities and the computer has only limited control over when the information it needs will become available. The challenge lay not so much in complexity--no deep and subtle reasoning was called for--but in the requirement for flexibility.

To no one's surprise, the systems that coped best with the spill problem were those designed as general purpose tools. These included OPS, RLL, AGE, ROSIE, and HEARSAY. For none, however, did the task prove easy, nor were the outcomes entirely satisfactory. Most of the teams had to build new features into their systems to handle dimensions of the problem that had not been encountered before. Because these features increased their systems' power and versatility, many planned to make them permanent components. And with tinkering, even the more narrowly focused systems proved adept at solving specific sub-tasks, like the source location problem. Plainly one goal of the conference had been met. The systems had been tested. Weaknesses and strengths, some anticipated, some surprising, had been revealed. As the week drew

to a close, the minor defeats and frustrations that everyone had experienced faded in importance. What remained was an exhilarating sense that the field's promise would be realized and that the day was not too far away.

The fate-tempting arrogance that the Greeks called hubris is said to be an occupational hazard among AI researchers. It seems to go with the territory. Peter Hart, formerly of SRI International, now with Fairchild Palo Alto Research Laboratories, captured its essence in verses, written at odd moments during the meeting, about a young researcher who "thought he could own/ all the wisdom that's known." Soon however, problems with a "subtle and curious" knowledge base and a sponsor furious about cost overruns brought him to his knees.

The point of this tale is not new,

But it may have escapeth a few:

Avoid, and don't bite

What, try as you might,

Exceedeth your power to chew.

DATE
ILME