computers and people

May — June, 1985
Vol. 34, Nos. 5-6

including *Computer Graphics and Art*

formerly *Computers and Automation*

SCIENCE FICTION WRITER USES PORTABLE COMPUTER IN HIS TRAVELS

The Computer Almanac and Computer Book of Lists – Instalment 41

Neil Macdonald
Assistant Editor

30 OF THE OVER 60 PRESENTATIONS AT THE ARTIFICIAL INTELLIGENCE AND ADVANCED COMPUTER TECHNOLOGY CONFERENCE AND EXHIBITION, LONG BEACH, CALIFORNIA, APRIL 30 TO MAY 2, 1985 (List 850501)

Text Generation Based on Mode of Comprehension / Ingrid Zuckerman and Judea Pearl / UCLA

Reasoning and Memory / Margot Flowers / UCLA

Graphical Mediators in Problem Solving / James Tanaka / Northrop Univ.

Real Time Contextual Analysis of Complex Sciences / Steven P. Smith, Ph.D. / Northrop Research & Technology Center

Automated Identification-Pattern Recognition & Learning Algorithms / Christopher Mayer / Consultant

Commercial Viability of AI in Medicine: How & Why / David J. Mishelevich, M.D., Ph.D. / Mishelevich Associates, Inc.

Verification of Medical Diagnoses Using a Microcomputer / Douglas D. Dankel II and Giuliano Russo / Univ. of Florida and Humana Hospital

A Paradigm for Real Time Inference / Robert C. Moore / LISP Machines, Inc.

A Prototype Expert System for Material Handling / Hatem N. Nasr / Univ. of Houston

Machine Remonitoring for the Factory of the Future / R. Gene Smiley and Richard L. Schiltz / Entek Scientific Corp.

Simulation of Paths Control for Automated Manufacturing / M. Luisa N. McAllister / Morovian College

Electrotopography Based Expert Systems and Machine Intelligence for Manufacturing in the Metallurgical Industries / M. Ensanian, B. N. Ensanian and T. A. Shaw / Electrotopograph Corp.

An Expert System for Space Shuttle Cabling / Roger Saxon and Roger Schultz / Abacus Programming Corp.

Engineering Applications of Expert Systems at Boeing / Janush S. Kowalik / Boeing Computer Services

AI-Based Technology for the Space Station / Oscar Firschein / SRI International

An Expert System for Voice Recognition / Russell B. Ives, / Univ. of Southern Calif.

Empirical Artificial Intelligence in Speech Recognition / Bill Meisel / Speech System, Inc.

Inside the Knowledge Workbench: A Prolog-Based Development Tool / Don Dwiggins / Silogic

Towards a Language for Knowledge Representation and Transformation / Thomas C. Brown / Kestrel Institute

Towards New Foundations of Intelligent Systems / Carl Hewitt / MIT

Architectural Classification for Expert System Design / N. S. Rajaram / Univ. of Houston

Robots: Can They Be Made Smarter? / Mysore Narayanan / Miami Univ. (Ohio)

Combining Demon-Based Parsing with a Phrasal Lexicon / Michael Dyer and Urni Zernik / UCLA

Evaluating Natural Language Systems: Tools vs. Solutions / Steven Shwartz / Cognitive Systems, Inc.

An Introduction to the Lingua Natural Language Parsing System / Kurt Fuqua / Manto, Inc.

Rule-Based Programming on Fifth Generation Computers / Gary Lindstrom / Univ. of Utah

(please turn to page 27)

Computing and Data Processing Newsletter

THE "ELECTRONIC COTTAGE": UPDATE

Based on a report in the "Boston Globe", Feb., 1985
135 Morrissey Blvd.
Boston, MA 02107

One of the more famous visions of how computers will define our future is a world composed of "electronic cottages." In this society, the huge organizational anthills, which today gather workers in by the thousands, will be replaced by computer networks that allow workers to co-ordinate their labor from their homes. To use the specific phrase, people will "telecommute" to work.

At the time these prophecies were first made the logic seemed inescapable. Employees would save money in transportation, food, child care, clothing, and by taking tax write-offs. They would benefit psychically by setting their own hours, working at their own pace, and going sailing when they chose. Employers would benefit through savings in rent and by tapping labor pools hitherto closed to them: the handicapped, mothers with young children, the semi-retired, prison inmates and persons in distant locations.

A Growing Network

In the years since these prophecies were first made the telecommuting infrastructure has grown more propitious every year. The microcomputers needed to participate in the networks have become relatively cheap and pervasive. E-mail and teleconferencing are now standard, well-understood, forms of communication. There is even an Electronic Cottager's Assn. that "meets" on CompuServe, one of the nation's information utilities, where professionals advise each other on the subtle points of rate structure and the mysteries of direct mail.

Yet, as inviting as the atmosphere has become, telecommuting has not caught on. The ratio of man-hours of work done at home to that done in the office may have shifted a little in the expected direction, but then again it may not have; figures on this issue are hard to gather. In any case, there certainly has been no drastic change.

Why is this? Many "futurists " attribute the lag to employer resistance. Some believe that managers fear losing control over their work force, and that until managerial techniques appropriate to the telecommuter are developed companies will discourage the idea.

But it is possible that the problem is deeper than that. In certain of its sections the Digital Equipment Corp. supports telecommuting as solidly as any company could: it gives workers the requisite hardware and pays the phone charges. John Redford works in one of these sections, and yet, he says, "true telecommuting is nonexistent. Everyone comes in at some time during the day."

Technology Lacking

The reason, he says, is that you just can't get enough information through the technology. The networks are slower than speech; you can't exchange drawings easily, and you and a colleague can't both go down the hall and examine some piece of equipment together. "In fact," he says, at least in his section, "if someone has not been (physically) in touch for two or three weeks, you can be sure that she or he has gone off in the wrong direction." Several important categories of information, including body language and voice tones, are difficult to convey over nets.

If this line of reasoning is right, then telecommuting does not make sense when the work being done is complex enough to require a fairly high level of communication. The paradox of this conclusion is that if the computer takes over the rote work and the mechanical procedures of the office, as it is supposed to do, then the work remaining for humans may actually demand a more intense level of interaction than before. If so, then the long-run impact of the computer may actually be to increase the proportion of work done in the office.

Ω

Vol. 34, Nos. 5-6
May-June, 1985

Editorial Offices	**Berkeley Enterprises, Inc.** 815 Washington St. Newtonville, MA 02160 (617) 332-5453
Advertising Contact	**The Publisher** **Berkeley Enterprises, Inc.** 815 Washington St. Newtonville, MA 02160 (617) 332-5453

"Computers and People" (ISSN 0361-1442), formerly "Computers and Automation," is published every two months at 815 Washington St., Newtonville, MA 02160, U.S.A., by Berkeley Enterprises, Inc. Printed in U.S.A. Second-class postage paid at Boston, MA and additional mailing points.

Subscription rates, effective Sept. 1, 1984: U.S.A., $18.50 for 1 year, $36.00 for 2 years; elsewhere, add $8.00 per year.

NOTE: The above rates do not include our publication, the "Computer Directory and Buyers' Guide." To receive this, please add $20.00 per year to your subscription rate in the U.S.A., and $23.00 per year elsewhere.

NOTE: No organization in Switzerland or Morocco is authorized or permitted by us to solicit for, or receive payment for, "Computers and People" or the "Computer Directory and Buyers' Guide." All subscriptions to and payment for these publications should be sent directly to Berkeley Enterprises, Inc.

Please address mail to: Berkeley Enterprises, Inc., 815 Washington St., Newtonville, MA 02160, U.S.A.

Postmaster: Please send all forms 3579 to Berkeley Enterprises, Inc., 815 Washington St., Newtonville, MA 02160, U.S.A.

Change of address: If your address changes, please send us both your new address and your old address (as it appears on the magazine address imprint), and allow four weeks for the change to be made.

computers and people formerly *Computers and Automation*
including *Computer Graphics and Art*

The magazine of the design, applications, and implications of information processing systems – and the pursuit of truth in input, output, and processing, for the benefit of people.

Editorial Note

We invite articles on the subject of computers and nuclear weapons. Computers, and computer people who work to make nuclear weapons work, are an essential ingredient of the nuclear evil.

There will be zero computer field and zero people if the nuclear holocaust and the nuclear winter occur. Every city in the United States and in the Soviet Union is a multiply computerized target. Thought, discussion, and action to prevent this holocaust is an ethical imperative.

Announcement

The Computer Directory and Buyers' Guide

Production of the master copy for printing of the 27th edition (1984-85) of *The Computer Directory and Buyers' Guide* has been further delayed. Meanwhile, any current subscriber to *Computers and People with Directory* who does not already have the last edition (26th, 1983) may on request to us receive a copy of that issue, so long as the overrun lasts.

Front Cover Picture

The front cover picture shows Arthur C. Clarke, science fiction writer, author of "2001: A Space Odyssey", and many other works, with his portable computer. The place is the Griffith Park Observatory in Los Angeles, CA. The nine pound microcomputer by Hewlett-Packard Company gives the author the freedom to write anywhere, at home or while travelling. Clarke's home is currently Sri Lanka (previously called Ceylon). One of his well remembered remarks of twenty years ago is, "I intend to go to the moon when tourist service starts."

Notice

"Computers and People" is not a technical journal and we do not supply free reprints or copies of articles. If a person wants an article, and cannot obtain the article from a nearby library, we can usually supply a back copy of the issue for $4.00 plus postage and handling. Also, microfilm copies of "Computers and People" (then named "Computers and Automation"), years 1969 and earlier, are available from University Microfilms, Inc., 300 North Zeeb Rd., Ann Arbor, Michigan 48106.

Key

[A]	–	**Article**
[C]	–	**Monthly Column**
[E]	–	**Editorial**
[EN]	–	**Editorial Note**
[F]	–	**Forum**
[FC]	–	**Front Cover**
[N]	–	**Newsletter**
[R]	–	**Reference**

Notice

*D ON YOUR ADDRESS IMPRINT MEANS THAT YOUR SUBSCRIPTION INCLUDES THE COMPUTER DIRECTORY. *N MEANS THAT YOUR PRESENT SUBSCRIPTION DOES NOT INCLUDE THE COMPUTER DIRECTORY

Editorial

Artificial Intelligence and Real Intelligence

Edmund C. Berkeley, Editor

Recently a good friend of mine wrote to me saying among other things:

> Your industry seems to continue to prosper, but I hope you never really succeed in your effort to substitute artificial intelligence for real intelligence.

This viewpoint raises several questions:

1. What is artificial intelligence?
2. What is real intelligence?
3. Will a machine ever be as intelligent as a human being?

"Artificial intelligence" has been defined as behavior by a machine which if performed by a human being would be called intelligent. The term was selected to sidestep the argument that only human beings could be intelligent.

There is no doubt that area after area of human activities that require intelligence are being performed by machines. The personal computer is an example. The expert system is another example. The autopilot in an airplane is still another example. Just as some tools apply mechanical power to tasks that human beings unaided cannot perform, so other tools apply electrical power to tasks that human beings unaided cannot perform.

What are those elements of behavior which persuade us to classify the behavior as intelligent?

Although one day it may be possible to take apart a human brain and observe how it actually handles information intelligently, yet nowadays the most we can do is to identify and illustrate some kinds of behavior that we classify as intelligent.

Let us consider some examples.

Addition. When you or I add 12 and 8 and make 20, we are behaving intelligently. We use our minds and our understanding to count 8 places forward from 12, for example, and finish with 20. If we could find a dog or a horse that could add numbers and tell answers, we would certainly say that the animal was intelligent. A computer of course can do this. And it can add more than 100,000 numbers in a second.

Selection. Or suppose that a dog or a horse going along a road comes to a fork and a signpost. If he could read the sign and then choose left or right depending on his destination and instructions, we would certainly agree that he would be acting intelligently. A computer can do this, more than 100,000 times a second.

Memory. The basic operation of intelligent behavior in the human brain is learning and remembering. Likewise, in a computer, the basic operation is storing information and referring to it. The number of locations in which information can be stored in a computer and its peripherals is more than 100 million. A human being from time to time may have to say "I have forgotten that," "I do not remember that." But a computer (like the proverbial elephant) never forgets.

Here is a list of dictionary definitions of intelligence:

- the ability to learn or understand from experience
- the ability to acquire and retain knowledge
- the ability to respond quickly and successfully to a new situation
- the ability to use reason in solving problems and directing conduct effectively
- the ability to deal with new situations
- the ability to deal with perplexing situations
- the ability to understand and infer in ordinary rational ways
- the ability to apply knowledge to manipulate one's environment
- the ability to acquire and apply knowledge
- the ability to think and reason

(please turn to page 27)

The "Strategic Computing Plan": An Assessment

Severo M. Ornstein
Brian C. Smith
Lucy A. Suchman
Computer Professionals for Social Responsibility
P.O. Box 717
Palo Alto, CA 94301

"Any computer system, however complex, and whether or not it incorporates AI, is limited in the scope of its actions and in the range of situations to which it can respond appropriately."

Reprinted with permission from the *Communications of the ACM*, February 1985, published by the ACM, 11 West 42 St., New York, NY 10036.

Computers: The Critical Field

In the 1940s, atomic physics was about 25 years old. Building on the discoveries of the new field, scientists were able to produce a weapon more powerful than had ever before been conceived. In the 1980s, computer science -- which also happens to be about 25 years old -- has become the critical field underlying modern weapon systems. This is not yet widely recognized; when we think of nuclear weapons, we tend to envision the warheads and the explosions, forgetting about the complex computer technology that supports the decision to fire the missiles and directs them to their targets. Computer systems are by now used throughout the military for early warning, communications, weapons guidance, and the simulations with which targets are selected and battles planned.

The "Strategic Computing Plan"

In the fall of 1983, the Defense Advanced Research Projects Agency (DARPA) issued a Strategic Computing Plan to develop a new generation of computing technology for military applications. The Plan initiates a five-year, $600,000,000 program, and there is good reason to believe this is just the beginning. The proposal contains plans for developing an underlying technology base of new hardware and software. The hardware emphasis will be on microelectronics and multiprocessor architectures, from which DARPA hopes to obtain at least a thousand-fold increase in net computing power. The software component focuses on artificial intelligence (AI) -- particularly on expert systems -- to provide machines with "humanlike, intelligent capabilities" /1/ including natural-language understanding, vision, speech, and various kinds of automated reasoning.

Completely Autonomous Air, Sea, and Land Vehicles

On top of this technology base, three specific military applications are to be developed. For the Army, the Plan proposes a class of "autonomous vehicles" able not only to move around independently but also to "sense and interpret their environment, plan and reason using sensed and other data, initiate actions to be taken, and communicate with humans or other systems." For the Air Force, the suggestion is a "pilot's associate" to aid aircraft operators who are "regularly overwhelmed by the quantity of incoming data and communications on which they must base life or death decisions," in tasks ranging from the routine to those that are "difficult or impossible for the operator altogether" and require the "ability to accept high-level goal statements or task descriptions." Finally, the Navy is offered a "battle management system" "capable of comprehending uncertain data to produce forecasts of likely events, drawing on previous human and machine experience to generate potential courses of action, evaluating these options, and explaining the supporting rationale." These three applications are intended to illustrate the power of the technology; we are also asked to imagine "completely autonomous land, sea, and air vehicles capable of complex, far-ranging reconnaissance and attack missions."

Automation of Military Decision-Making

Two facts stand out. First, the Strategic Computing Plan proposes the use of AI technology in military systems in order to provide a radically new kind of flexibility and adaptiveness. Referring repeatedly to the increased speed and unpredictability of modern warfare, the Plan promises that computing technology can be developed that is capable of adapting to "unanticipated enemy behavior in the field." /2/ This will require "a new generation of military systems" that could "fundamentally change the nature of future conflicts." The change involves both increasing the amount of computation and enlarging its role to include automation of military decision making.

Close Military Control of Research

Second, there are specific proposals about how to direct computer science research. Rather than letting researchers follow their own course, the Plan aims to focus them on military objectives. Various mechanisms are suggested to do this, such as close coupling of fundable research goals and military needs, adherence to strict developmental timetables, and the selection of specific development projects intended to "pull the technology-generation process." (The Army, Navy, and Air Force projects cited above are the first examples.)

Hoping for Reliable Decision Making When It Is Impossible

In assessing the Strategic Computing Plan, our concern is not with the underlying technology base, or with military projects as such. Nor do we question the power of AI as a new and important technology. Our concern is that increased reliance on automated decision making in critical military situations, rather than bringing greater security, leads in an extremely dangerous direction. In suggesting such a role for AI, the Strategic Computing Plan creates a false sense of security in the minds of both policymakers and the public. The problem is that the Plan hopes for reliable decision making in circumstances where there may simply be no way to achieve it -- with computers or with humans.

Three Interacting Trends

Modern warfare is marked by three interacting trends: increasingly powerful weapons, more separation (in both time and space) between planning and execution, and a faster and faster pace. The first means that the consequences of our actions, intended or unintended, can be greater that ever before. The second means that we rely on increasingly large, complex, and indirect systems for command, control, and communication. The third means that any miscalculation can quickly lead to massive ramifications that are difficult, perhaps impossible, to control. It is easy to see the dangerous potential of the three in combination. They are all the direct product of technological developments in offensive and defensive weapons systems; and they have brought us to the situation that we live with now: two nations confronting each other with forces that, if unleashed, would destroy both in less than an hour.

The Current State Is Precarious

This danger is recognized on all sides; people differ only in what they think we can or should do about it. However, if anything is universally accepted, it is that the current state is precarious; and into this situation the Strategic Computing Plan proposes to introduce AI as a new ingredient:

> Improvements in the speed and range of weapons have increased the rate at which battles unfold, resulting in a proliferation of computers to aid in information flow and decision making at all levels of military organization. ... A countervailing effect on this trend is the rapidly decreasing predictability of military situations, which makes computers with inflexible logic of limited value. ... Confronted with such situations, leaders and planners will ... be forced to rely solely on their people to respond in unpredictable situations. Revolutionary improvements in computing technology are required to provide more capable machine assistance in such unanticipated combat situations. ... Improvements can result only if future computers can provide a new "quantum" level of functional capabilities. [pp. 3-5]

Unpredictability Will Be Very Great

What this means in plain English is the following: Faster battles push us to rely more on computers, but current computers cannot handle the increased uncertainty and complexity. This means that we have to rely on people. However, without computer assistance, people cannot cope with the complexity and unpredictability, either. So we need new, more powerful computer systems.

The role that computers are to play, furthermore, is not minor; the Plan makes clear that reliance on automatic systems is meant to include the control of strategic weapons. For example,

> Commanders remain particularly concerned about the role that autonomous systems would play during the transition from peace to hostilities when rules of engagement may be altered quickly. An extremely stressing example of such a case is the projected defense against strategic nuclear missiles, where systems must react so rapidly that it is likely that almost complete reliance will have to be placed on automated systems. At the same time, the complexity and unpredictability of factors affecting decisions will be very great. [p. 4]

People With Judgment and Common Sense Needed to Cancel False Alarms

The Strategic Computing Plan offers no argument to warrant this reliance on automatic decision making. Any computer system, however complex, and whether or not it incorporates AI, is limited in the scope of its actions and in the range of situations to which it can respond appropriately. The ballistic missile warning systems of the United States, for example (and presumably those of the Soviet Union), are designed to err on the side of oversensitivity, and regularly give false alarms of incoming attacks. /3/ Although most of these alarms are handled routinely, on a number of occasions they have triggered the early stages of a full-scale alert. These false alarms stem from causes as varied as software inadequacies in dealing with natural events (in one case a moonrise, in another a flock of geese), failures in the underlying hardware (such as a failing integrated circuit chip that started sputtering numbers into a message about how many missiles were coming over the horizon), and human error (such as that of an operator who mounted a training tape onto the wrong tape drive, thereby causing the system to react seriously to what was intended to be a simulation). The primary insurance against accidents resulting from this kind of failure has been the involvement of people with judgement and common sense. So far, there has always been enough time for them to intervene and prevent an irretrievable, and perfectly real, "counterattack."

Unanticipated Events

Despite these lessons, the Strategic Computing Plan promotes the view that the human element in critical decision making could be largely, if not totally, replaced by machines. This would require that computers embody not only "expert knowledge" but also common sense and practical reasoning. What distinguishes common-sense reasoning, however, is the ability to draw on an enormous background of experience in the most unpredictable ways. In directing a friend to your house, for example, you do not have to give instructions about all the possible things that might happen along the way: fallen trees, accidents, flat tires. An extraordinary range of knowledge and experience may be relevant; we never know what we will need or when we will need it. Nor do we usually even notice that we are using this background knowledge. These facts undermine any attempts to codify common-sense knowledge and practical reasoning. As a result, current expert systems do not have the common sense of even a small child. This lack of common sense means that in AI systems, as in any computer system, unanticipated events are liable to trigger anomalous reactions. This is particularly a problem with military systems since, as the Strategic Computing Plan points out, it is the unpredictability of war that poses the gravest threat.

Many Unfulfilled Promises From Artificial Intelligence

Sophisticated AI systems are scientifically intriguing; they enable us to explore areas of human capability in which we have enormous interest, including those areas that are relevant to coping with uncertainty. Over the years, the lure of AI has led to a growing appetite for research funding. This appetite in turn has led the professional community to make promises, many of which have turned out to be more difficult to fulfill than was anticipated. For example, it was widely believed in the 1950s that we would soon have fully automatic machine translation, an accomplishment that still eludes us. These unfulfilled promises are frequently a combination of ordinary naivete, unwarranted optimism, and a common if regrettable tendency to exaggerate in scientific proposals. Shortcomings are often masked by subtle semantic shifts. When we fail to instill "reasoning" or "understanding" in our machines, we tend to adjust the meaning of these terms to describe what we have in fact accomplished. In the process, we obscure the real meaning of our claims for the power of AI.

Unrealistic Confidence in the Power of the Technology

When claims are taken literally, without appropriate qualification, they give rise to unrealistic confidence in the power of the technology. Policymakers, even those close to the profession, are not immune to such misconceptions. Witness the following discussion of Defense Department research on space-based weapon systems, as reported in the "Los Angeles Times" on April 26, 1984:

> ... The fireworks began when a panel that included Robert S. Cooper, director of the Defense Advanced Research Projects Agency, George Keyworth, Reagan's science adviser, and Lt. Gen. James A. Abrahamson, director of the Strategic Defense Initiative, acknowledged that a space-based laser system designed to cripple Soviet long-range missiles in their "boost" phase would have to be triggered on extraordinarily short notice.
>
> To strike the boosters before they deployed their warheads in space would require action so fast that it might preclude a decision being made in the White House -- and might even necessitate a decision by computer, the panel said.
>
> At that, Sen. Paul E. Tsongas (D.-Mass.) exploded: "Perhaps we should run R2-D2 for President in the 1990s. At least he'd be on line all the time."
>
> "Has anyone told the President that he's out of the decision-making process?" Tsongas demanded.
>
> "I certainly haven't," Keyworth said.
>
> Sen. Joseph R. Biden, Jr. (D.-Del.) pressed the issue over whether an error might provoke the Soviets to launch a real attack. "Let's assume the President himself were to make a mistake ... ," he said.
>
> "Why?" interrupted Cooper, "We might have the technology so he couldn't make a mistake."
>
> "OK," said Biden. "You've convinced me. You've convinced me that I don't want you running this program."

A Profoundly Human Political Problem

Cooper's final comment betrays the belief that computers could be competent to take over critical decisions and might correct deficiencies in human judgment as well. As the discussion shows, common sense suggests that such a claim is implausible. It might have been that common sense was wrong -- that the underlying science had advanced beyond the layperson's expectations. However, we believe that the skepticism is in fact well founded.

To cope with problems of complexity and speed in modern warfare, the Strategic Computing Plan proposes a quantum leap in computer technology comparable to the advent of nuclear-weapons technology in the 1940s. Ironically, the problems arise in part from the very technology that is proposed as a solution. Past attempts to achieve military superiority by developing new technology, rather than increasing our security, have brought us to the present untenable situation. The push to develop so-called "intelligent" weapons as a way out of that situation is another futile attempt to find a technological solution for what is, and will remain, a profoundly human political problem.

References

/1/ Unless otherwise noted, all quotations are from Strategic computing, new-generation computing technology: A strategic plan for its development and application to critical problems in defense, Advance Research Projects Agency, Rosslyn, Va., Oct. 28, 1983.

/2/ Schwartz, L. DoD.to get $95 million in funding. "Electron. News 3", 1489 ar. (Mar. 19, 1984), 18.

/3/ See, for example, the Hart-Goldwater report to the Committee on Armed Services of the United States Senate, Recent false alerts from the nation's missile attack warning system, U.S. Government Printing Office, Washington, D.C., Oct. 9, 1980, and Accidental nuclear war, "Newsl. Physicians Soc. Responsibility 3, 4" (Winter 1983), 1.

Ω

The Joy of Computer Chess

David Levy
International Master
c/o Prentice-Hall, Inc.
Englewood Cliffs, NJ 07632

"One of the most controversial aspects of computer chess is the question of whether or not a computer program can eventually be stronger than the (human) World Chess Champion."

Consisting of selections from *The Joy of Computer Chess* by David Levy, published by Prentice-Hall, Inc., Englewood Cliffs, NJ 07632, 1984, 129 pp, and reprinted with permission.

Editorial Note: *This is a fascinating book for any person who plays the royal game of chess, and for any person who can put a chess program on his or her microcomputer. — ECB*

Table of Contents

Preface

This book is for chess enthusiasts, computer enthusiasts and, above all, for computer chess enthusiasts. The chess player will enjoy playing over the games, and he will learn a lot of detail about how computer programs play chess, how they 'think', what are their strengths and weaknesses, and how to use and enjoy a chess playing program to best advantage. The computer enthusiast who is interested in writing a program for his own computer will find sufficient information in this book to enable him to perform such a task from start to finish, and sufficient ideas to provide endless hours of pleasurable experiment. The computer chess enthusiast will enjoy the book for all sorts of different reasons.

Within a few years there will be a personal computer and/or a chess enthusiast in almost every home in the civilized world, and the subject matter of this book will then have become of universal interest. I have tried to make the text easy to understand, both for chess enthusiasts who have no computer knowledge whatsoever and for computer literates who know no more than the rules of chess. My purpose in writing this book has been to bring the joy of computer chess to the millions of people who realise what a wonderful game chess is, and to the millions who have their own computer or who are thinking of buying one. I very much hope that I will succeed in this aim. ....

Chap. 1 Position Representation and Move Generation

The first problem to overcome when planning to program a computer to play chess, is how to tell the computer what chess is. It is one thing for a human to gaze at a chessboard, see where the pieces are located and understand the relationships between the various pieces, but a computer is merely a device that can store and manipulate numbers. So how, exactly, do we teach a computer what chess is all about?

The human recognizes the chess pieces by their shape and size. The pawns are the smallest and the kings and queens the largest. One player's pieces are one colour, his opponent's pieces are another colour. All of this must be conveyed to the computer so that it can understand where the pieces are on the board and can calculate which moves the pieces may make. This is accomplished by assigning a different number to each piece type.

A simple scheme might assign the number 1 to a pawn, 2 to a knight, 3 to a bishop, 4 to a rook, 5 to a queen and 6 to a king. A white piece can be designated by a positive number, a black piece by a negative number, and an empty square by zero. Using this method of piece representation, the initial position in a game of chess would look like this:

Black

-4	-2	-3	-5	-6	-3	-2	-4
-1	-1	-1	-1	-1	-1	-1	-1
0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0
1	1	1	1	1	1	1	1
4	2	3	5	6	3	2	4

White

It is not difficult to appreciate, when studying the above diagram, that the program can always detect whether a particular square is occupied, and if so by what piece (and of which colour), simply by testing the number stored in the relevant location.

The above representation is the simplest possible, but it fails to convey everything that we might wish to know about the position. ....

Chap. 5 Best Computer Games

In this chapter there are a number of games which have been played by computer programs. Some of them were played in computer tournaments, and so both players are programs. Some of the games were played against strong human players. Some of the games were played by programs running on 'mainframes' (big computers costing millions of pounds), others were played on microcomputer systems costing only a few hundred. I have selected games which will show the reader just what computer programs are capable of, and in some cases, where the games are 5 years old or more at the time of writing, they show that for some years now computers have had the ability to produce surprising results against very strong human players.

In the case of games played at blitz speed, the rules are a little different to those used in human v human blitz games. The human player has 5 minutes in which to make all of his moves, but because of the overhead in move transmission time and in having a human move its pieces, the program is usually allowed to make 60 moves, at an average of 5 seconds of processor time per move. If he reaches move 60 without being mated, the human wins on time.

White: CHESS 4.6
Black: Michael Stean
(International Grandmaster)
Blitz Game. London, September 1977
Owen's Defence

((At this point the next 22 moves on each side are given, with comments, and then there is the comment from a player:))

"Bloody iron monster" exclaimed Stean, who only now realized that his queen is needed to prevent mate.

((Moves 23 to 27 are then given, and:))

"This computer is a genius," said Stean.

((Then moves 28 to 39 are given, and Stean resigned.))

((The rest of the chapter consists of reports on 7 more games in which one or both of the players is a computer program.))

Chap. 6 How Strong Can Computers Become?

One of the most controversial aspects of computer chess is the question of whether or not a computer program can eventually be stronger than the (human) World Chess Champion. Some experts in Artificial Intelligence predicted as long ago as 1968 that it was only a matter of time, and not very much time at that. In fact two of the world's leading authorities in the field, Donald Michie and John McCarthy, started a bet with me in 1968 that I would lose a match against a computer program within 10 years. The bet grew in size, the programs got stronger, I got weaker, but I still won the bet in September 1978. As the period of that bet drew to a close, and it was clear that I was favourite to win, one of the world's leading experts on computer chess, Monty Newborn, wagered a few hundred dollars that by 1984 there would not be a human player who could stand up to the best computer programs. Furthermore, as recently as 1979 Newborn went into print with the prediction that '... it is highly probable that programs will be playing Master level chess by 1984, Grandmaster by 1988, and better than any human by 1992. (These are conservative estimates!)'

After I collected my winnings (except those from a certain person who welched), there were others ready to start a new bet, and I have a $1,000 wager with former President of the Association for Computing Machinery, Dan MacCracken, that I can survive a match against the strongest program at the start of 1984. These bets show that a number of highly intelligent people, who are amongst the world's leading experts in various branches of computer science, are confident that Master, Grandmaster and even World Champion level chess are all within the scope of the computers of tomorrow. The only question is, when will tomorrow come?

No-one knows the answer to this question. My own opinion is that it will be between the years 1995 and 2000 when a computer program can play at the level of a strong Grandmaster under tournament conditions. My reasons for this estimate are intuitive, and clearly it is difficult to establish a scientific method of estimating when a program will play as well as Bobby Fischer, but let us see what the trends have been during the past decade or so.

In one sense at least, chess programs can already be said to be intelligent. The famous British mathematician Alan Turing devised a test for intelligence in computers, whereby a human is allowed to interrogate a program via a teletype. If the human is unable, after 'conversing' for some time, to tell whether the teletype is hooked up to another human being or to a computer, then if it is a computer at the other end of the line, then that computer can be said to be intelligent. The best chess programs have already passed the Turing Test in chess, since they can play well enough to prevent a strong human player from knowing what is at the other end of the line. To test this hypothesis, a West German television station conducted an interesting experiment during a simultaneous exhibition being given in Hamburg by International Grandmaster Helmut Pfleger. Three of the humans who were 'playing' in the simul, were sitting with small earphones in their ears. In the balcony of the playing hall other humans were watching their games through binoculars, and feeding Pfleger's moves into computer programs. When the programs responded, their moves were relayed to the appropriate human via the earphones.

Pfleger suspected nothing, and even after the exhibition was over and he was told that three of his opponents had been computer programs, he found it difficult to believe, especially so when he learned that one of the programs, Ken Thompson's BELLE, had been his successful opponent in a well played skirmish. So BELLE, at least, has passed the Turing Test. .... Ω

Berkeley – *Continued from page 20*

ly not important, and base our conduct on the real and huge accumulations of differences that need to shake and alter our basic assumptions. The megaton nuclear bombs exploded by both Americans and Russians, the ballistic missiles created by both, the earth satellites orbiting around the earth, the computing mechanisms that launch and guide them, make a real and huge accumulation of differences.

References

/1/ From "Tyranny on Trial: The Evidence at Nuremberg", by Whitney R. Harris. Reprinted by permission of the Southern Methodist University Press, Dallas, Texas, 1954.

/2/ From "Applied Logic", by W.W. Little. Reprinted by permission of Houghton Mifflin Co., Boston, Mass., 1955.

Ω

Walking Machines — Part 2

Prof. Robert McGhee
Director, Digital Systems Laboratory
Ohio State University
1659 North High St.
Columbus, OH 43210

"The area we are working on most intensively at the present time is terrain adaptive vehicles, . . . roughly speaking, an artificial horse."

Excerpted from *Intelligent Systems,* edited by Jean Hayes and Donald Michie. Distributed by Halstead Press, a division of John Wiley and Sons, Inc., New York. Reprinted by permission of the publisher. Part 1 of this article appeared in the March-April 1985 issue.

A Quadruped Vehicle

The General Electric Quadruped Vehicle was constructed during the beginning of 1964, completed and first tested about 1967. The machine uses the human being for both the logical and co-ordination levels of control. He is given a power assist and does not have to provide more than 1% of the full force associated with limb motions. Thus, when the man moves his legs or his hands the machine amplifies that motion and applies 100 times more force than the man is applying. So the man is not receiving very much assistance here except for power amplification. He must himself co-ordinate 12 joints, three in each leg, two at the hip and one at the knee. Despite this very heavy load on the operator (twice as many degrees of freedom as a helicopter), a few men were able to achieve a remarkable degree of dexterity in operating this machine. The principal difficulty was that the work was exhausting. Even the best operator, who happened to be the designer of the machine, was only able to operate it for about 10 minutes a day: then he was exhausted. When I became aware of this machine, I thought that we ought to be able to convince the people concerned that they needed an auto-pilot. So I decided with one of my students, who is now a professor, Dr. Andrew Frank, to build the machine which was subsequently christened The Phony Pony. It had only one purpose -- to show that joint motions could be co-ordinated electronically. We had no microprocessors at that time so we built our own special-purpose computer. It had just 16 flip-flops in it and it still took a rack of equipment. But it proved the point. We showed that motion could be co-ordinated electronically, and did not necessarily require biological intelligence.

I showed a short film about these ideas at a meeting in Yugoslavia in 1969. It attracted more attention from researchers in the Soviet Union than in the United States. They went back home and began to work themselves on the problem of motion co-ordination, and by 1972 had solved a major scientific problem. The scientific problem which was solved was: How should one co-ordinate the motion of the limbs of a walking machine in order to maximize its stability? The solution is what is called a wave gait. It involves a wave of motion from the rear of the machine to the front of the machine, with the placing of legs on the righthand side and the lefthand side half a cycle out of phase. Through a very intensive, long-term computer study, the Soviets proved that this was the optimal way to use the legs of a six-legged machine. The extra pair of legs had been added because it has been found that for machine co-ordination of motion, trying to produce an artificial quadruped was too ambitious. The balancing problem is too difficult. Thus the Soviet walking machine has an extra pair of legs in the middle to simplify the stabilization problem.

This produces 18 degrees of freedom, three in each leg (that is the minimal number needed to place each leg arbitrarily). With so many degrees of freedom (18 compared to the six or seven in an industrial robot) a very rich kind of behaviour is possible. There is not a unique trajectory associated with the body motion over a given terrain but one can decide what to control. You can decide, for example, that it's important to keep the body level. And that can be done. You can decide that, unlike a wheeled or tracked vehicle, it is not necessary to bounce over every stone but do what a human being or a horse does, either stepping over it or on it. And in this way, rather large obstacles can be overcome.

An Animal Adjusts Its Body

I saw a film of this work in 1972 and was very impressed. It is quite important to realize that what was shown was not an artist's conception, but an allegedly real-time film in which a computer was making all of the decisions concerning motion. However, I came to the conclusion that this kind of idealized motion was not what was really wanted. At least it is not what animals do. Animals adjust their body so that it is more or less parallel to the local terrain slope. This is to avoid two kinds of problems: it is unpleasant for an animal to scrape its belly and so it likes to keep it off the ground; on the other hand, if it gets its body too high, it runs out of leg length and that is equally embarrassing. So it seems more sensible to follow the biological solution in most circumstances and not try to keep the body level but rather roughly parallel to the terrain.

By 1976 the Soviets had made another advance, having become aware of the importance of some kind of remote sensing. To be effective, a robot (Soviet work has been concentrated on true robots, machines without human operators on board) needs some kind of primitive vision. It needs to discover obstacles before it runs into them.

So they did an excellent simulation study in 1976 in which the action of a triangulation range-finder is assumed. This is not imaging vision. It is felt in the Soviet Union, and I agree, that at this point imaging vision systems are too difficult and too advanced for control of locomotion. But we can have something that amounts to radar: a device that scans a terrain and provides a relief map, using the mechanism of triangulation. The scanners are on top of the vehicle, the receivers are in the body and if you compute the line of intersection of the image with the source, you then know where the terrain is. This is still to date the most advanced demonstration I am aware of, of a walking robot finding its way through difficult terrain.

Also in 1976 construction was under way both at my institution in Ohio State and at several points in the Soviet Union, to realize in hardware the kind of behavior demonstrated by simulation studies. Our machine, called the Ohio State University Hexapod, first walked early in 1977, and to the best of my knowledge, the first Soviet machine walked about four months later. Since 1977 we have been studying the problem of higher level control: How can a man and a computing machine co-operate effectively to regulate and co-ordinate the motions of the joints of the robot?

Communication Between Operator and Walking Machine

We recently made a film to demonstrate two kinds of communication between the human operator and the robot. The first kind of communication is symbolic, using a greatly simplified language. The language consists simply of isolated commands of which there are about a dozen. One is "Try out your legs and see if they are working OK". Thus, if any portions of the system are not functioning we discover this before trying to walk. Before walking it has to take its mark, positioning its limbs properly for the initiation of walking. For the second type of communication with the robot we need something like a pilot's controls. This particular scheme involves a three-axis joystick. Twisting the joystick causes the robot to twist its body -- determines its rotational rate, in fact. Fore/aft deflection determines fore/aft velocity, right/left deflection determines right/left velocity and of course co-ordinated motion is possible. So now the operator only thinks about what he wants the body of the machine to do and the computer worries about what the legs should do.

When the robot walks, the motion is quite slow. This is because of another defect in this machine that we have discovered, through experience. I referred earlier to the stubbornness of a mule. A mule won't harm itself. In fact, it is more intelligent than a horse. You can force a horse to harm itself by asking it to do things that exceed its muscular capacity. A mule won't do it.

This machine is even more agreeable than a horse. It will do whatever you tell it to do -- such as ripping its front legs off. We think it important that a real walking machine, a practical walking machine, must be given an artificial sense of pain so that it will refuse to do things that will damage the machine or harm the operator. Currently, machine capabilities include co-ordinated motion, forward motion, turning, and some motion to the side. It is capable of side-passing, i.e., it can keep its body in a constant orientation and move strictly laterally. To back up, one simply pulls the control level to the rear and if it is twisted at the same time then the vehicle begins to turn.

The computer is not on board, though there is a considerable amount of on-board digital electronics. The computer provides via cable three pieces of information for each leg: namely, the position of each of the three joints, the velocity of each of the three joints, and (at present for one leg) the force of reaction with the ground, so the computer knows what forces the robot is encountering as it goes over the terrain. This feature permits terrain adaptation and efficient use of energy. Before long, we expect to be able to put railroad ties (sleepers) around the floor of the room and repeat this experiment with the machine passing over the railroad ties and making the adjustments automatically.

For this it has to have a vertical sense, which it now has, with a vertical gyroscope and two pendulum sensors. When the experiment is over we again ask it to stand up.

Japanese Quadruped Machine

The above is a short account of recent work in Ohio State University. We are in contact not only with researchers in the Soviet Union but also in Japan, where some very advanced work is going on. In particular, my own early ideas about logical or finite state control have been picked up by Japanese researchers, Professor Hirose in particular, and advanced to a considerably higher level to produce a quadruped machine with several interesting features. First of all, the legs were designed by a mechanical engineer (I am an electrical engineer). A mechanical engineer has sense enough to get the actuators off the limbs where they simply produce unnecessary weight and consequent energy inefficiency. He put the motors in the body and moved the joints through tendons, much like the human body. Another feature of the machine is that it has a tactile sense which allows it to discover the presence of obstacles and adjust to them. A defect of the machine is that it is too autonomous. There is no human communication. When it is sent out on a path it finds its own way and you cannot call it back. Another defect is that having only four legs, it has an inferior degree of stability. It is not able to move its legs and its body at the same time, which illustrates again the necessity for more than four legs.

The first commercially viable machine that I would be willing to call a walking machine is the Menzi Muck mentioned earlier. It is very clever and apparently cost-effective, selling for $80,000 to $100,000. In one case it has been adapted to timber harvesting. The hydraulic shear on the arm is able to cut a tree of up to 12 inches at one snip. And the operator can pile the trees up for later collection and processing.

The OSU Hexapod walking machine is something analogous, perhaps, to a walking stick insect which is not very advanced. If you look at more advanced insects, there is limb specialization. The grasshopper is an example. Its back legs are used for jumping, its front legs are used for holding food and its inner legs are for stabilization. The same kind of specialization is seen in the Menzi Muck because this is a machine for a specific function, not a research machine. The co-ordination is very primitive, however. I think that productivity could be improved if the operator had a microprocessor assisting him. He ought only to have to concern himself with what the end effector is doing and not pull levers constantly to move joints one by one. When he wants to move to a new site, the power comes from the arm. The rear legs -- their motion is simplified by the use of wheels -- and the inner legs are used for support and stabilization.

Present Status

That is where we are today. Within three years we hope to have developed a new man-carrying machine derived from the OSU Hexapod vehicle. It will probably have an internal combustion engine like a Toyota with about 80 horsepower. The machine is likely to weigh 3,000 or 4,000 pounds and will certainly at the beginning have six legs. We would like to be able to get rid of the middle pair of legs but do not think we can deal with the stabilization problem adequately without them.

The actuation mechanism is likely to be hydraulic, although that is a serious problem. The efficiency of conventional hydraulic actuators is unacceptable. We expect the operator will have aircraft-type controls and will be concerned with controlling speed and direction. On-board computers, at least

(please turn to page 28)

The Social Responsibility of Computer Scientists

Edmund C. Berkeley
Berkeley Enterprises, Inc.
815 Washington St.
Newtonville, MA 02160

The logical fallacy is the failure to respond realistically to a real change when a real change has occurred.

Based on Chapter 16 of *The Computer Revolution* by Edmund C. Berkeley, published by Doubleday & Co., Inc., Garden City, NY, 1962, 249 pp, reprinted with permission.

Are Computer Scientists To Be Judges of Social Responsibilities?

The essential question, the essential argument, that needs to be discussed and analyzed is this one:

> The social responsibility of computer scientists is a topic in the field of ethics, in the field of the social sciences, not a topic in the field of computers and data processing. A computer scientist cannot be expected to be competent as a social scientist. He is hired to do a job; he is not hired to think about the consequences or implications of his work with computers. This is outside of his territory of competence, and is the concern of his employer.
>
> Computer scientists have no special social responsibilities as computer scientists, only the responsibilities of all scientists and citizens.

There are a number of rebuttals to this argument, though when one suddenly encounters it for the first time, one may not be able to answer it well and clearly.

To make the answer as vivid as possible, let us begin with the story of a certain locksmith.

The Story of the Locksmith

Once there was a man who was in the business of making locks and keys, and who was very skillful. One day a stranger walked into his shop and said to him, "I want you to make a key which will open a certain safe." The locksmith said to him, "Whose safe is it?" The stranger said, "Never you mind whose safe it is. I will pay you handsomely for the key. I'll blindfold you, and take you to the place where the safe is. You can have all the tools you want -- I'll pay for them -- and you make me a key. Besides, while you make the key, you will have a chance to work out some intensely interesting scientific theories, and after the safe is cracked open, I will give you permission to publish some papers, those that don't reveal too much information. Think it over, I'll be back tomorrow."

So the locksmith wondered about the remark "Never you mind," and the blindfolding, and the secrecy; but he knew it was hard enough to earn a living, and the promises of the stranger sounded attractive and exciting. So he said to himself, "Well, that fellow would just get another locksmith if I did not go," and so he decided he would go. And the next morning the stranger came for him, and he allowed himself to be blindfolded and went.

For several years the locksmith tried to open the safe, and then at last he succeeded. But the stranger did not allow him to look inside; all the locksmith saw was the door swing open. The stranger then said to him, "Here is your pay -- now go away -- and remenber not to talk about this -- or you will get into a lot of trouble."

After a few more weeks, the locksmith read in the newspaper that what the stranger had taken out of the safe was a supremely intelligent directing mechanism for flying weapons, from the size of a wasp to the size of an eagle, which would enable him to pinpoint and exterminate any person, any community, any town, any city in the whole world. And he read the stranger's declaration that henceforth the world was to do exactly as he commanded, and that any opposition to his commands or dictates would be precisely and completely destroyed.

The Questions Presented

This story presents us with four questions at least: Is the story entirely fictitious and impossible? Was the stranger a criminal? Could the locksmith have recognized the stranger as a criminal? Did the locksmith do what was right?

The story, of course, is more parable than it is fiction. We know with sadness the many points where it agrees with the facts of past and current history, and predictions of the future.

The Criminality of the Stranger

As to the second question, it seems to me that it is not necessary for us to argue the criminality of the stranger because that has been settled, by the Nuremberg trial after World War II of the leaders of Nazi Germany. This trial is reported well in "Tyranny on Trial: The Evidence at Nuremberg" by Whitney R. Harris. The book has an introduction by Justice Robert H. Jackson of the Supreme Court of the United States, who participated in that trial; and contains a full story of the trial of the German war criminals. The author, Harris, served as trial counsel on Justice Jackson's staff. The report is an extraordinary, breath-taking, and bloodcurdling story, worth careful reading to show how and in what way the German state under Hitler planned, prepared, and carried out aggressive war under a thick screen of lies. Let me now quote from Chapter 38, "The Law and Aggressive War" (p. 514ff.): /1/

> ... In the first few years of the thermonuclear age there has been placed in the hands of men a new power potential capable of such destructiveness as to threaten the users of the power as well as the intended victims. War has always been homicidal; now it has become suicidal. Civilization may see an end to war, because it cannot survive a renewal of war. The second factor [possibly causing the end of war] is the universal condemnation of aggressive war, of which the Nuremberg judgment is both source and reflection. For many years prior to World War II, the peoples of the world had thought of agressive war as wrongful and wicked. The Nuremberg judgment gave expression to that feeling by punishing the individuals responsible for launching World War II.

> ... The difficulty of applying the concept [of aggressive war] in close cases does not mean that courts are powerless to recognize inexcusable aggressive action when it clearly occurs.

> ... The defendants could not have been surprised as to the moral aspects of their conduct. No one sends millions to die without a qualm of conscience.

> ... Aggressive war does not become defensive war by the simple act of calling it such.

> ... The slaughter of civilians in concentration camps, ordered by Hitler, was a crime of Hitler, even though he directed his mass killings as head of the German state.

> ... It is after all moral condemnation which underlies legal prosecution. The killing of innocent human beings by order of heads of states is subject to substantially the same moral blame whether it is the killing of civilian populations in connection with war or the killing of troops resisting unlawful aggression.

> ... Of course, no one should be heard to assert absolute immunity for acting in accordance with the orders of anyone else, even in such a fundamental matter as war.

At almost the end of the chapter, the International Military Tribunal is quoted:

> "War is essentially an evil thing. Its consequences are not confined to belligerent States alone, but affect the whole world. To initiate a war of aggression, therefore, is not only an international crime; it is the supreme international crime differing only from other war crimes in that it contains within itself the accumulated evil of the whole."

> Harris continues:

> This statement is law, and what is more, "This law applies for all times, in all places, and for everyone, victor and vanquished." The initiating and waging of aggressive war is now indisputably criminal. No more important decision was ever made by any court.

It seems to me that this settles the second question, the criminality of the stranger; it settles the law and the morality; the wrongfulness and wickedness of aggressive war; of sending millions to die with or without qualms of conscience; the moral condemnation that underlies legal prosecution; and the impermissibility of arguing immunity for acting in accordance with the orders of someone else.

Recognition of the Stranger As a Criminal

As to the last two questions, the responsibility of the locksmith for recognizing the stranger as a criminal and for doing what was right, there is no doubt that according to law a locksmith has to satisfy himself that a customer has a bona fide right to the locksmith's help in opening a safe. Locks and keys and safes have been in existence long enough for the judgment of society to agree that a locksmith must satisfy himself that a man who comes to him to open a safe has a good right to have the safe opened. The more valuable the goods in the safe, the more necessary is the examination of the stranger, and the more important is the responsibility to do what is right.

The Towering Problem of Our Time

So much for the general argument. Now for the specific example, the towering problem of our time, intercontinental ballistic missiles with megaton nuclear warheads guided by computing mechanisms. In this case, three groups of scientists play the role of locksmith: the men who make the nuclear warheads, who are the atomic scientists; the men who make the rocket motors that will propel the missiles; and the men who make the guidance systems, the computer scientists. Let us talk about the computer scientist.

The computer scientist, according to law and morality, does not have the right to shut his eyes in regard to the stranger, no more than the locksmith has. Both have to keep their eyes open.

The computer scientist like the locksmith must judge the stranger. The stranger will not say what his real purpose is. The stranger, in fact, may be altogether unable to say what his purpose is; he may be in the grip of strong psychological forces (paranoia, for example) that he has no understanding of whatever. Certainly Hitler did not consider that he himself was a psychopath. But deeds speak louder than words, and the locksmith must look at the deeds.

Therefore, let us set up a number of criteria for the locksmith to decide what is the purpose of the preparations for opening the safe. For example, in the case of an arms race between two countries A and B, in order to decide what these preparations really mean, the locksmith can make up a long list of objective tests:

Test 1: Does country A have armed bases surrounding country B? and vice versa?

Test 2: Is country A (or country B) increasing or decreasing its military forces? expanding or contracting its testing of nuclear weapons? ...

Test 3: What are the claims announced by each country for political or territorial changes, which probably can be obtained only by force?

Test 4: Can the economy of country A (and country B) remain stable and function well without heavy war preparations?

The computer scientist, like the locksmith, has the moral and legal duty to study these questions and answer them objectively. He does have a special responsibility because without him the safe cannot be opened.

And so we arrive at the proposition that it seems to me we have to support:

> Computer scientists have a special responsibility as computer scientists, more than and in addition to the responsibilities of most other scientists and citizens -- the responsibility of the locksmith.

The Avoidance of a Certain Logical Fallacy

In all the discussion and argument about the social responsibility of computer scientists, there is undoubtedly present in the minds of all of us a logical fallacy that is begging continually to be accepted as truth, because it is truth so much of the time. We want to go on with "business as usual". We do not want to see that "something new has been added." We want to work tomorrow and the next day on our usual problems and not think about the new ones. We do not want to act on the basis that something has really changed, that a new and most terrible power has come into the possession of the governments of the United States and the Soviet Union.

In the possession of the United States and the possession of the Russians are more than enough nuclear explosives to put an end to the life of man on earth. We want to say "yes, that may be, but somebody will do something about it, and I do not have to make any change in what I am doing."

The logical fallacy is the failure to respond realistically to a real change when a real change has occurred.

The same fallacy, refusal to realize a change, operated widely before World War II. In 1938 the government of Great Britain under Chamberlain and the government of France under Daladier, renounced their agreement with Czechoslovakia, and told Hitler that it was all right for him to take the Sudetenland from Czechoslovakia, since Hitler said this was his last demand; and Chamberlain returned from Munich announcing "Peace in our time!" The governments of Great Britain and France were in the grip of this same fallacy, the failure to see that a new and very real and terrible change had actually occurred.

In one of the best books on logic that I know, "Applied Logic" by W. W. Little, W. H. Wilson, and W. E. Moore, this fallacy, the failure to acknowledge a real change, is given a name, "The Argument of the Beard." From the book: /2/

> In a sense, the argument of the beard may be considered the opposite of the black-or-white fallacy. We are guilty of the black-or-white fallacy if we fail to admit the possibility of middle ground between two extremes. We are guilty of the argument of the beard if we use the middle ground, or the fact of continuous and gradual shading, to raise doubt about the existence of real differences between such opposites as strong and weak, good and bad, and white and black.
>
> ... The fact that we cannot determine the exact point at which white ceases to be white does not prove that there is no difference between white and black.
>
> The very name of the fallacy is derived from the difficulty of deciding just how many whiskers it takes to make a beard. Surely one whisker is not sufficient. Possibly even 25 are too few. Then let us say that 350 whiskers make a beard. Why not 349? 348? and so on. We would have trouble determining an exact minimum. Does this fact mean that there is no difference between having a beard and not having one? ... If a car can carry seven persons in an emergency, why not just one more? By the argument of the beard, a car should be able to carry an infinite number of passengers. ... This error ... is especially pernicious in value judgments because it is frequently used to justify unethical conduct.

Avoiding the Fallacy

We may guard against the argument of the beard by reminding ourselves that although a difference may be small, it may nevertheless be so real that an accumulation of such differences may bridge the distance between great extremes.

Now, a computer scientist may say:

> "Well, it makes no difference if I work on an early-warning radar network, and the computers that go with it, because I help to defend my country against attack."

And another computer scientist may say:

> "I am not doing any worse than that fellow in the early-warning radar network, because I am working on the guidance system for an air-to-air missile which can be used to knock down an enemy missile to be detected by the early-warning system."

And a third computer scientist may say:

> "Well, it is true that I am working on the guidance system for an intercontinental ballistic missile, BUT it will only be launched if an enemy ICBM comes over to destroy one of the cities in my country."

And a fourth computer scientist may say:

> "Well, I am working on the computations relating to the spread of poison gases, BUT I am very sure that my country will only use poison gas if the enemy uses poison gas."

And finally some kind of mistake occurs in the whole tragic pattern -- information comes in that poison gas has been used when in fact it has not, or that ICBM's are on the way when in fact it is only the moon rising over the horizon, or some poorly maintained computer in a distant country has a failure. ... Then all these fine differences count for nothing at all -- ICBM's land on Moscow, Leningrad, and Kiev, others land on New York, Chicago, and Los Angeles, and at least 40 million human beings are dead -- more deaths in a day from less than twenty bombs than all the deaths of World War II in six years from all weapons combined.

Inattention to the accumulation of small differences constitutes the fallacy of the argument of the beard, the fallacy of failing to acknowledge a real and huge change because it has been gradual.

We need to stop justifying our conduct on the basis of small differences that are real-

(please turn to page 13)

Towards Machines That Think

N. N. Sachitanand
c/o Kasturi & Sons, Ltd.
Kasturi Bldgs.
Madras 600002, India

"Heuristics is better known as 'rules of thumb' and has also been called the 'art of good guessing'."

Based on two articles in *The Hindu* Dec. 4 and 5, 1984, published by Kasturi & Sons, Ltd., Madras, India and reprinted with permission.

Fantasy Becoming Fact

Machines that can think like humans are no strangers to science fiction. Feature films of recent times have also featured "intelligent" machines like HAL of "2001 -- A Space Odyssey" or R-2 D-2 of "Star Wars". However, in the last 30 years a new area of research has been gathering momentum in the real world which promises to turn this fantasy into fact at the turn of the century. This field of research is now well known as "Artificial Intelligence," a term coined at a meeting in 1956 at Dartmouth College, U.S. One of the primary subjects on the agenda of that meeting, mooted by Marvin Minsky (now at MIT), John McCarthy (now at Stanford), Nathaniel Rochester of IBM and Claude Shannon of Bell Laboratories, was to discuss ways of simulating thought with computers.

Intelligent Machines

One of the earliest attempts at devising a machine to emulate human thought was made by Charles Babbage in Britain in the mid-nineteenth century. However, his mechanical computer, called the Analytical Engine, which was to be programmed by punched cards, remained an unfinished invention. With the development of the digital electronic computer during the Second World War and its subsequent versions of today with their fast throughput capabilities, it seemed that intelligent machines would soon be in the market place.

But when AI researchers got down to devising "intelligent" systems using current computer technology and techniques, they soon learnt that human thought processes are not easy to simulate. True, fundamentally, many of the activities of the computer, as we know it today, resemble human-cognitive processes. Both mind and machine accept information, manipulate symbols and store/retrieve items from memory. But the essential difference lies in how the human and the computer arrive at a conclusion.

Human Reasoning by Analogy

Present day computers, even the giant number crunchers, are fantastic devices for running through bits of information at high speeds. But they do so serially and to reach a conclusion have to go through an exhaustive and very explicit sequence of steps written in the program. In attempting to develop machines that mimic human thought, AI investigators have found that humans rarely reach conclusions in a sequential manner.

In an article by Patrick Huyghe in "Psychology Today," Donald Norman, Director of the Institute for Cognitive Science at the University of California (San Diego), is quoted as saying: "Humans reason more by example. We rarely think or make decisions by standard rules of logic. Human reasoning seems to operate more by means of analogy and experience."

According to John Anderson, Professor of Psychology and Computer Science at Carnegie-Mellon University in Pittsburgh (one of the three major centers for AI research in the U.S. along with MIT and Stanford University), "a major component of intelligence is the ability to convert much of one's knowledge into knowledge that can be used in pattern matching. These patterns or combinations of information allow the expert mind to quickly recognize situations and deal with them according to experience."

Heuristics

he human expert seems to arrive at his ision not so much by deductive reasoning by means of a set of rules of judgment d logic, derived from experience, applied to a collection of facts. These judgmental rules are referred to by AI researchers as heuristics, a word with the same Greek root as eureka. Heuristics is better known as "rules of thumb" and has also been called the "art of good guessing". It enables an expert to recognize promising approaches to problems, to break problems down into smaller problems, to get around incomplete data and to make educated guesses when necessary.

Expert Systems

The difference between the working method of a normal computer programmed conventionally with rigid algorithms and a human expert using heuristics can best be illustrated by considering the game of chess. Today's biggest computer working full time would need a decade to consider all the potential moves early in a game and arrive at the best move. A human expert on the other hand, would consider only the most likely moves, weigh their potential based on his storehouse of chess knowledge and determine which move to make, all within the time allotted in the rules of the game.

Recognition of the above fact has helped AI researchers develop the first breed of marketable "intelligent" machines called expert systems which use a combination of a knowledge base and heuristics to solve problems in limited areas such as determination of the structure of complex molecules or diagnosis of faults in locomotives and the like. The heuristic rules come from discussions between an expert in the field and a computer-wise person trained to translate the results of the expert's out-loud thinking into a computer program. This technique for translating an expert's experience into symbols that computers can understand has come to be called "knowledge engineering".

New Computer Languages

To use heuristics, programmers have had to develop whole new families of computer languages since currently used computer languages like Fortran and Pascal are designed for manipulating numerical problems and not well suited to AI applications. One of the most commonly used AI languages is Lisp, for List Processing. It operates by linking lists of data. It can match lists, concentrate lists, shuffle them, take them apart or do whatever is needed to get the information desired. Dialects of Lisp exist, such as Q Lisp, which operates on sets of symbols rather than lists of things. Most expert systems in the U.S. employ Lisp. In Europe and Japan, a language called Prolog is favored. A Prolog program starts with a logical statement and tries to determine if it is true or false. A statement might be "London is north of Rome". The knowledge base may contain facts such as "London is north of Paris" and "Paris is north of Rome". When found, these would prove the truth of the original statement.

Knowledge Bases

Expert systems have necessarily to use very large knowledge bases. To manage such big knowledge bases efficiently, many expert systems use what are called "if -- then" induction statements. When specified "ifs" are satisfied, they lead to "thens" that represent new concepts or solutions to problems. For example, "if" a person has a running nose, a temperature and is sneezing, "then" the person has a cold.

Besides "if -- then" rules, knowledge can also be represented as tree-shaped networks of related objects or concepts, known as semantic networks. In these networks, facts are clumped together in "nodes" which are interconnected by all possible paths such as "is -- a" (Example: A sparrow is a bird) or "has -- a" (Example: Birds have wings). These pathways, when followed, lead to a conclusion or inference (Example: A sparrow has wings).

A recent article in "Technology Review" by Joel N. Shurkin relates how the first program to use heuristics emerged from a conversation at Stanford between AI researcher Edward Feigenbaum and Nobel Prize-winning geneticist Joshua Lederburg. Feigenbaum wanted to see if he could emulate in a computer the kind of empirical deduction common to the scientific process. Lederburg suggested beginning with the analysis of organic compounds using mass spectroscopy. Lederburg immersed himself in computer science and also recruited Carl Djerassi, a respected professor of chemistry at Stanford.

The task was stupendous. They had to determine the basic concepts involved and develop rules that express the relationship between concepts. A major part of the challenge was for the computer specialists to find out what Lederburg and Djerassi knew and how they knew it and then convert that information into symbols that a computer could understand -- a give and take process that took several years.

Chemists know that the structure of any chemical compound depends on a number of basic rules about how atoms bond to each other. When they make or discover a new compound, chemists can analyze the substance with a mass spectrograph, which provides a lot of data about the compound but no clues as to the specific shape that the molecule takes out of the millions of possible shapes allowed by the rules of chemical bonds.

Building the right kind of "if -- then" program to narrow the range of possibilities was only the first major problem to be solved. The knowledge engineers had to dig out from the chemists how the latter determined molecular structure from the spectra and then add these judgmental rules to the program.

First Commercial System: Dendral

Eventually, in 1965, the interdisciplinary team came up with an expert system called Dendral which was able to predict a narrow range of possible structures of a compound from its spectral data. Dendral was the world's first commercial expert system and is now regularly used by organic chemists. It has been further expanded to include data from other analytical techniques such as nuclear magnetic resonance.

Mycin

In the mid-1970s, Edward Shortliffe, a Harvard-educated premed found himself at Stanford studying computer science and medicine. His adviser was Stanley Cohen, soon to be famous for his work on recombinant DNA. Shortliffe produced a program -- incorporating the expertise of Cohen and physician Stanton Axline -- that would diagnose blood and meningitis infections and advise physicians on antibiotic therapies. The program, called Mycin, performed at the level of human specialists in infectious diseases and above the level of general physicians. By adding another program called Teirasias, the system was also able to tell the consulting physician the reasons for the guesses it made.

Mycin had other problems that limited its usefulness in clinical situations, but it served well as a lesson in knowledge engineering. In developing Mycin, researchers found that if they removed the knowledge base from the program -- that is, the medical information -- what they had left was a section that contained the logic and this section was universally applicable. You could plug in data bases from other fields, say geology or computer-chip design, and the program would still work.

Inference Engine, Emycin

Researchers now call this logic portion the "inference engine" and they have developed a program incorporating this generalized logic named Essential Mycin or Emycin.

IBM is now using an expert system based on Emycin to diagnose malfunctions in computer disk drives and Sacon is another Emycin-driven system that assists structural engineers in identifying the best strategy for using a complex computer simulation program. The ability of the inference engine to operate in several fields supports the notion that at least some human reasoning is structural and can be duplicated by machine.

New Ventures

The development of the expert system as a practicable proposition has touched off a mad rush in the U.S. to move AI from the laboratory into the marketplace. Like another futuristic technology, genetic engineering, the investment is being made at both ends of the specturm -- small companies started by AI researchers themselves backed by venture capital and large corporate giants like ITT, General Electric, DEC and Texas Instruments.

"Business Week" in a recent cover story reports that venture capitalists have injected more than $100 million into some 40 small companies bent on commercializing AI. Startups in this field tend towards science fiction names like Teknowledge (started by a group that included Feigenbaum in 1981), Machine Intelligence Corporation, Computer Thought Corporation, Symbolics and Intelligenetics.

Optimistic analysts predict that the market for expert systems and the software tools needed to build them will explode from about $20 million this year to nearly $2.5 billion by 1993. Simultaneously, a market is developing for software that allows users to communicate with expert systems in normal or "natural" languages. It is expected to grow to $1.8 billion annually by 1993.

Assortment of Strategies

About 200 researchers in the U.S. -- 500 worldwide -- are working on developing expert systems and there are about 50 systems in a stage of readiness. A number are now commercially available but only half-a-dozen actually make money for their developers. This is because expert systems are very expensive to develop. An expert system can take many man-years of work by knowledge engineers, who are scarce, and can eat up $1 million or more. Worse, the chances of its performing well are difficult to predict. Companies in the business have devised an assortment of strategies to make their technologies attractive. A Californian start-up company offers low-priced systems for use on personal computers, some have made a speciality of teaching knowledge engineering to their customers' programmers and others build custom-tailored systems for clients.

Other vendors and would-be users pin their hopes on the recent emergence of an industry in expert system "shells". These are off-the-shelf inference engines that users can equip with special expertise for doing anything from analyzing the course of pollution in streams to selecting the right dinner wine. Taylor Instruments of Rochester, New York, uses "shells" from several suppliers to build systems for controlling processing plants.

Expert System Applications

Some companies have developed systems in-house which may have a broader market and become a lucrative sideline. For example, General Electric has developed a program that helps mechanics repair ailing diesel engines on its locomotives. Although the system was originally designed for GE mechanics, the company now plans to offer it to other railroads. GE intends to develop similar expert systems for jet engines and digital flight control systems.

Expert systems are today being used or being designed for use in such disparate fields as medical diagnosis, insurance underwriting, credit planning in banks, prospecting for minerals, production scheduling, equipment maintenance, battlefield intelligence, data analysis, spare parts ordering, product design and what have you. "Expert systems can be consulted for advice," says Randall Davis of MIT, "or they can be co-workers on a more equal footing. Or they could be assistants to an expert. All along the spectrum these are very useful systems."

While expert systems seem to be on a good commercial wicket and find plenty of real world applications, their capabilities just barely touch the fringe of what human intelligence can do. An article in "Fortune" by Tom Alexander points out that the main connection between expert systems and intelligence is that most systems employ programming techniques pioneered by AI researchers trying to develop computer models of how the mind works.

Even prominent AI researchers like Marvin Minsky of MIT and Roger Schank of Yale contend that today's expert systems are largely based on 20 year-old programming techniques that have merely become more practical as computer power got cheaper.

"Brittle" Systems

Expert programs use languages that manipulate non-numerical symbols and emulate the deductive operations of classical logic. But the computer still works mechanically, manipulating arid symbols that it recognizes but does not understand. It does not deal in rich associations, metaphors and generalizations that language evokes in people and that constitute the essence of meaning and thought.

The rule-bound expert system could break down if it encountered a situation it was not programmed for. Not knowing what it didn't know, it would be bound to provide misleading responses while having no idea it had done so. For example, several expert systems can diagnose automobile breakdowns. They know all about ignition and fuel problems and might be a terrific help to novice mechanics. But they would be frustrating time wasters if the problem is, say, a stone in the exhaust pipe. The term that AI people use for this shortcoming is "brittleness".

"If a Rule Is Wrong, the Expert System Won't Tell"

Beau Sheil, a manager of Artificial Intelligence systems with Xerox observes: "It's extremely difficult to clone experts except in very, very narrow, highly specified domains. A limit of expert systems is that if the rules are wrong or don't apply in this case who tells? One thing that is sure is that the expert system doesn't tell".

Sheil concludes that the main applications of expert systems will be to help specialists manage information or point out inconsistencies between their assumptions and established knowledge. Martin Hollander, director of marketing with Intellicorp, sees a system functioning best as a "doubting Thomas" adviser. "You give the system a hypothesis," he says, "and the system runs through its rules and points out things it knows are inconsistent".

Common Sense

Most humans, in contrast, have the advantage of common sense. They can handle the unusual by making analogies with their experiences, by extrapolating missing information or by modifying imperfect instructions. One of the objectives of frontline AI research today is to attempt to program common sense into a computer. This is a stupendous task since the knowledge base will require acquisition and representation software powerful enough to handle millions of facts and tens of thousands of rules. To get that knowledge into a computer, AI researchers believe that machines must be able to learn on their own.

How Humans Learn

Researchers have gone back and studied how humans learn. Models of children learning geometry problems and language reveal that we do not simply soak up knowledge like sponges. "What has emerged", says psychologist John Anderson, "is basically a set of principles of how we learn from doing. It turns out that we learn formal skills like solving problems in physics and doing proofs in geometry not by reading a textbook and understanding the abstract principles, but by actually solving problems in those fields. What the textbooks don't teach you is when to apply the knowledge and that knowing turns out to be about three-quarters of the learning problem."

According to a theory of semantic memory developed several years ago by M. Ross Quillian at Carnegie-Mellon, the mind is an enormously complicated and constantly changing network of nodes and links. When we experience something new, like seeing an exotic animal, we store the information and later retrieve it through a technique called "spreading activation". This means that new material is processed by being linked with established ideas or modes. That way, the new animal is not only classified according to form, color, odor and behavior but also linked to other animals and a repertoire of feelings and recollections.

This rich network of connections in human memory is one of the most profound differences between humans and machines. The brain's ability to search for information through its millions of neurons simultaneously looks positively uncanny. "Spreading activation makes use of associations among ideas," Anderson says, "so that when a conversation turns to restaurants, for instance, all the knowledge related to the subject becomes instantly available." Quite unlike the computer the more information we have about a subject, the faster we seem to be able to retrieve it.

"Learning by Discovery"

Yet, human memory is less than ideal at times. Information sometimes gets lost, because we break up experience into bits and pieces and store them in different parts of memory, according to Roger Schank, a professor of computer science and psychology at Yale. But on the positive side, he notes that breaking up of knowledge in memory allows us to make better generalizations and more useful predictions.

Douglas B. Lenat, a Stanford University computer scientist, believes that a shortcut to machine learning exists in what he calls "learning by discovery". In this method, the computer automatically acquires both new knowledge and the rules by which that knowledge is handled. To achieve it, Lenat loads a knowledge base with an initial set of symbolically expressed concepts. These are manipulated by rules that combine the concepts in various ways, then evaluate the results. Many combinations yield nonsense but a few produce new concepts.

Analogies and Metaphors

Lenat developed an expert system called Artificial Mathematician, which employed 100 concepts of set theory and 250 rules. By acting on the concepts, the rules discovered natural numbers, prime numbers, addition, subtraction, multiplication and division. If it is possible to discover new concepts in this way, Lenat reasoned, then it should be possible to discover new heuristics, since the rules themselves are concepts. He went on to develop an expert system called Eurisko, which generates and evaluates new heuristics in a variety of areas. Lenat hopes to use Eurisko to represent knowledge in such powerful cognitive mechanisms as analogies and metaphors, which are the basis for common sense.

The huge knowledge bases needed for common sense reasoning beget another problem:

finding efficient ways to quickly retrieve and update information. The trouble with present computer designs is that operations are performed sequentially. A machine checks one "if" at a time, in an "if-then" system to determine whether it applies to the problem at hand. When a knowledge base grows to tens of thousands of rules, sequential search becomes so slow that direct human interaction with the program becomes impracticable. Present-day solutions involve algorithms that group rules so that only potentially pertinent groups are searched. But future solutions will entail new computer architectures that will permit parallel or simultaneous data retrieval and data processing.

Nearly 50 universities in the U.S. are currently working on so-called parallel processing machines in which a vast number of micro-computer chips (each with a microprocessor and its own memory bank) are linked together to perform separate parts of a job. At the University of Maryland researchers have constructed a system called Zmob which uses 256 parallel microprocessors connected to a host minicomputer and arranged in a high speed circular conveyor belt that allows for simultaneous handling of messages. At Columbia University a machine called Dado is planned which will have 1,023 microprocessors to search different parts of a rule base simultaneously. MIT is developing the Connection Machine which will have a mind-boggling one million microprocessors wired to each other.

The task is not easy. Processors, like people working together, need to be synchronized when searching for data or solving problems. They also need to communicate, so that two of them are not trying to do the same operation.

Parallel Programs

Parallel machines do not make it any easier to write better AI programs. Good models for parallel programs and code languages in which to express them do not exist. The greatest problem involves the absence of algorithms for breaking problems into independent pieces that can be processed separately and then put back together.

AI researchers are anxiously awaiting the so-called Fifth Generation computers which will use multiple processors to gain faster access to their huge memory banks and perform tasks at much faster speeds than the current breed of super fast computers. In an article in "Electrotechnology", W. S. E. Mitchell mentions that where current computers can perform 1,000 to 10,000 of the inferences used in expert systems per second, the Japanese are planning to achieve in their Fifth Generation computers, to be ready by the 1990s, a performance of 100 million to 1 billion inferences per second.

Another goal is to build a dataflow machine consisting of from 1,000 to 10,000 microprocessors, storage of one to 10 Gigabytes and a speed of from 1 billion to 10 billion instructions per second. It is with such computers equipped with interfaces that can accept and answer queries in natural language that AI systems will be able to truly match human intelligence.

Knowledge Acquisition

Another problem faced in developing AI systems is the transfer of expertise from the specialist to the computer program. "This is the critical bottleneck in AI," says Stanford's Feigenbaum. "It is the greatest research problem that AI laboratories must face and solve in the coming decade."

The present give and take interaction between the expert and the programmer is very time consuming. Although a knowledge engineer can get the basic information from an expert and other sources like textbooks in a week or so, it can take many worker-years to refine the program. Stanford's Bruce Buchanan is working on a concept called "knowledge acquisition" in which the expert would talk directly to the computer without having to go through a programmer.

Creativity

When one hears about all that AI has accomplished or is going to accomplish, the questions that naturally arise are: Will computers be creative? Will they be able to feel emotions? Marvin Minsky says there is no substantial difference between ordinary thought and creative thought. He believes that we take ordinary thinking so much for granted that we never wonder how it happens until a particularly unusual performance attracts attention. Then we call it genius or creativity.

What actually seems to separate the ordinary thinker from the extraordinary thinker is that the latter has learnt to be better at learning. If that is all there is to it, Minsky says, then once we can get machines to learn -- and learn to learn better -- one day we might see creativity happening in ma-

Sachitanand – *Continued from page 26*

chines. In fact, an international prize amounting to a handsome sum is being offered to the first AI program which comes up with a mathematical concept hitherto unknown to man.

Extension of Brainpower

Once ordinary human thinking has been programmed into a machine, Minsky believes that even emotions will be programmable. "It is a mistaken idea that feeling and emotion are deep whereas intelligence and how we get ideas are easy to understand," he says. However, as Professor Raj Reddy of Carnegie-Mellon put it in a talk at Bangalore, what is important is not that AI techniques can make computers fly into a rage or weep but that they provide the means for man to extend his brainpower extragenetically, just as he has extended his physical prowess by artifacts.

"What you see now is the next stage of human evolution," he says. Edward Feigenbaum also sees the coming of a new age. "We humans are very good at converting sensory signals to cognitive signals and at solving problems that require common sense. But in the face of large amounts of data we quail: we are unsystematic and forgetful, grow bored and get distracted. Writing and book technology helped us overcome some of these problems; interactive smart computers will help some more. We should give ourselves credit for having the intelligence to recognize our limitations and for inventing a technology to compensate for them."

Ω

CACBOL – *Continued from page 2*

A Distributed Implementation of Functional Program Evaluation / Joe Fasel / Los Alamos National Lab.

Intelligent Job Aids: The Use of Small Expert Systems in Business / Phil Harmon / Consultant

Expert Systems with Intangible Output / James H. Johnson / Human Edge Software

Extension of the Relational Data Model for Knowledge Base Management / David Hartzband / Digital Equipment Corp.

(Source: announcement of conference by Tower Conference Management Co., 331 W. Wesley St., Wheaton, IL 60187)

Editorial – *Continued from page 6*

As we read over this list, we observe that not a single one of these abilities is beyond the power of a computer appropriately programmed. In fact, if we consider the context of game-playing, the world champion in checkers is a computer program, the world champion in backgammon is a computer program, and the best computer chess programs today play better than 95 percent of human chess players who are members of chess clubs.

So the answers to our questions are:

1. Artificial intelligence is remarkably real intelligence.
2. Real intelligence includes thousands of instances of artificial intelligence.
3. Machines are already more intelligent than great numbers of human beings.
4. But human beings have great capacities of intelligence that machines do not yet have, and may not have for centuries.

Ω

10 APHORISMS (List 850502)

Incompetents often employ capable assistants.

Don't ask a barber whether you need a shave.

If the facts don't fit the theory, it may well be that the facts should be steamrollered.

If something seems to be doubtful, it is possible to make it sound convincing.

It is not wise to count your songbirds, until each one sings.

If it looks easy and quick, it is difficult and slow. If it looks difficult and slow, it is close to impossible.

Some ideas cannot be thought by the human mind.

Fortunate is he who expects nothing, for he shall not be disappointed.

No experiment is ever a complete failure, if it is looked at from a sufficiently broad point of view.

Nothing will be attempted if all possible objections must first be overcome.

(Source: Neil Macdonald's notes.)

Ω

McGhee – *Continued from page 16*

one microprocessor per leg, at least one central co-ordination computer, and probably a completely independent fault detection and fault correction computer will be included to provide safety for the men and the machine. The whole thing will, we hope, travel at something like 5-8 miles per hour. The speed may not be impressive, but we hope its mobility will be.

I believe that we understand the co-ordination problem quite thoroughly. On the other hand we do not know exactly how decision-making should be apportioned between the man and machine. That is, we do not understand the logical level as thoroughly as we understand the geometric or kinetic levels.

But there is another problem which is really serious. There is a curve derived from work done 30 years ago by Gabrielli and von Karman and shows that for all known modes of transportation there is a kind of an asymptote along that line which relates top speed to specific power. The vertical axis shows the size of the engine per ton required in a vehicle, beginning with 0.1 horsepower per ton and going up to 2000 horsepower per ton. As the speed goes up, of course the power requirement goes up. And it is interesting to note that (to no-one's surprise) horses are very efficient. A horse can move a ton with a little less than one horsepower or a little more than a horsepower, depending upon how fast he is moving. He is much more efficient than a human being. If you want to get around efficiently, apparently it is better to use four legs than two. Human beings require roughly three or four times more energy than horses. They are not able to go as fast and are not able to travel as efficiently. We are still very much better than tracked vehicles for off-road locomotion. Tracked vehicles require roughly 10 to 20 times more horsepower per ton than a horse. A tank requires several hundred horsepower and weighs tens of tons. Walking machines are still further up. Despite their adaptability, the energy costs are unacceptable. So a major challenge to us, and half of our program at my university, involves moving the walking machines' curve downward on the graph. The mechanical engineers think they know how to do it.

Ω

#20