# COMPONENT PART NOTICE (CON'T)

## COMPONENT PART NOTICE

THIS PAPER IS A COMPONENT PART OF THE FOLLOWING COMPILATION REPORT:

(TITLE): Artificial Intelligence in Maintenance: Proceedings of the Joint Services
Workshop Held at Boulder, Colorado on 4-6 October 1983,

(SOURCE): Denver Research Inst., Colorado.

SEP 1 7 1984

To order the complete COMPILATION report use __AD-A145 349__.

THE COMPONENT PART IS PROVIDED HERE TO ALLOW USERS ACCESS TO INDIVIDUALLY
AUTHORED SECTIONS OF PROCEEDINGS, ANNALS, SYMPOSIA, ETC. HOWEVER, THE
COMPONENT SHOULD BE CONSIDERED WITHIN THE CONTEXT OF THE OVERALL COMPILATION
REPORT AND NOT AS A STAND-ALONE TECHNICAL REPORT.

THE FOLLOWING COMPONENT PART NUMBERS COMPRISE THE COMPILATION REPORT:

# Artificial Intelligence Contributions to Training and Maintenance

Joseph Psotka
U.S. Army Research Institute

## Abstract

Artificial intelligence is rapidly becoming a practical and useful technology for training and maintenance. This paper provides an introduction to its uses in maintenance training, drawing on current research funded by the Army. After a description of this work, a call is made to fund more exploratory research, expand the base of competent professionals in the field, and begin the complicated process of evaluating this new technology in order to diagnose its failings and hasten its development.

## Introduction

My sincere thanks and appreciation go to our hosts for this workshop on artificial intelligence and maintenance training. The timing could not be more appropriate as the very many fine activities described in this volume testify. The development of training systems, particularly maintenance systems, is about to be complicated in quite new and, in many respects, unanticipated ways. Artificial intelligence (AI), knowledge-based systems, and expert systems for training and maintenance are becoming an overnight success after some 20 years of hard effort. We must prepare ourselves to use them wisely. The complicated issues that have been raised in connection with computer-based instruction in the past (cf. Orlansky, 1983) are all reissued in new guise: there is, after all, a continuity between this new application of AI and the older (but still developing) forms of CBI and simulations. We must be careful. Yet, the promise is also so clear that we must be prepared to act quickly and not squander a very real opportunity.

### What is Artificial Intelligence?

The phrases "artificial intelligence," "knowledge-based systems," and "expert systems" are used interconnectedly throughout this paper. None of them have precisely defined meanings. AI is the most generic of the three, since it involves activities like pattern recognition and voice synthesis and recognition that are generally not central to knowledge-based systems and expert systems. Expert systems are the mind of any AI program. An expert system incorporates inside a computer all the rules of action or thought a human expert has about any well-defined domain of knowledge or skill. For instance, John McDermott's Vax configuration program incorporates 500 rules extracted from a professional who was paid very highly for just knowing these rules. It must have been very disappointing to find out that a lifetime of expertise could be summed up in so few rules.

Knowledge-based systems are all expert systems, and they are the best known and understood kind of expert system. Very little is known about how to use expert system technology for training highly practiced skills and automatic skills like driving tanks and shooting, but arcade and video game technology is rapidly pushing in this direction.

All of these AI applications employ the symbol manipulation powers of the new generations of computers to solve problems. Problem solving seems to be at the heart of all expert systems. Expert systems technology provides a way to extract knowledge from an expert and place it into computer symbolic forms (in a process called knowledge engineering) and then use that knowledge base to solve problems. Problem solving in this context has a very broad definition, covering things ranging over algebra problems, medical diagnosis, mechanical or electrical maintenance and troubleshooting, all the way to training field generalship or even running an automated tactical operations center.

## Opportunities for the Army

Three rapidly evolving developments make the time ripe for applying artificial intelligence approaches to the Army's severe training and instruction problems. The major development is rapidly improving hardware that has produced personal LISP machines at low cost (20 - 30K) and the power of last year's mainframes. The second development is in the application of "expert systems" technology to a host of real world problems (ranging from medical diagnosis to mineral prospecting) that have demonstrated the utility of artificial intelligence techniques in very dramatic style. Finally, current state-of-the-art in computer- based training and instruction has advanced to the stage where the leap to using "expert systems" approaches is practical and possible: the basic knowledge is there to model soldier, task, and instructor characteristics with fidelity in a machine.

These technological windows of opportunity are opening in time to address some pressing Army needs and challenges. The Army of the 1990s is increasingly a high-technology Army and this is imposing tougher new demands on its soldiers. Maintenance training needs to become much more sophisticated and individualized. The distributed battlefield of the future will also make unprecedented demands on the cognitive decision making skills of its soldiers. They need to be prepared intellectually to make fast, appropriate decisions and use complicated strategies and technologies in order to win the war. What better way exists than to use high technology in order to train soldiers effectively to use high technology on the battlefield?

## Previous Work

During the past several years, ARI has engaged in a concerted program of research and development to prepare for this activity. Work on computer-based instructional systems has progressed on several fronts, ranging from the hand-held, talking, technical term tutor to a computerized videodisc

306

instructional system (SDMS). Training and simulation work has progressed from 3-D computerized systems like AMTESS to the current proposals for AI-based maintenance tutors. The groundwork for specific artificial intelligence applications has been carefully laid with a 1M per year jointly funded ARI/ONR program supporting basic research in artificial intelligence applications to maintenance and training. This work has reached the advanced development stage where several projects can lead to working demonstrations in the next few years if the critical funding is made available.

## Current AI Applications

The Army is in the middle of several initiatives that have the potential of providing significant funding to AI research and development in general, and maintenance training research in particular. Let me try to provide a brief overview of these programs.

- ARI will provide funding to several AI activities, to be described in more detail later in this paper.

- ARO is in the process of selecting a U.S. institution of higher learning with graduate level programs to develop a capability for research, education, and training in AI to help fill Army needs.

- TRADOC has formed a General Officer Steering Committee to oversee the development of an overall Army AI action plan by November 30, 1983.

- DCSRDA in concert with several other Army agencies, including DARCOM, HEL, TACOM, ETL, and ARI, has developed specific demonstrator projects, including:

      1.    Reconnaissance Vehicle;
      2.    Ammunition Supply Robot;
      3.    Intelligent Display;
      4.    Expert Medical System;
      5.    Intelligent Maintenance Tutor.

All of these activities are providing high level and very visible support to AI research and development within the Army. It is certain that these plans will amount to something significant, so we can look forward to a period of productive research and development.

## Current Academic and Commercial Efforts

Although there are plans for applying AI techniques to Army problems, we must examine the larger world of academic and commercial development for examples of how these systems work and what training purposes they might serve.

The best known and most mature developments of expert systems for problem solving have occurred in medicine: CADUCEUS by Myers and Pople at CMU, and MYCIN by Shortliffe at Stanford. Both systems are designed for decision making and diagnosing diseases to help an expert physician interpret a case. This is an important point because it has led people to believe that expert systems can only help experts of a similar kind (in this case doctors) do their job better. That is basically a faulty conclusion: expert systems can both act as consultants to experts, and they can replace experts. However, most of us do not yet have to worry: there are some rather severe restrictions on the kinds of expertise that can be entered into expert systems. Medicine provides virtually a prototype for this kind of expertise.

1. The knowledge domain must be bounded and well defined. It must be small enough to be manageable; yet large enough so that conventional, linear algorithms and programs will not work.

2. The problem or task domain is intellectual rather than dependent on sensory or motor skills. The problems are defined by a large number of interacting variables that require human knowledge to resolve.

3. At least one and preferably more human experts exist who can solve the problems. These experts must be able to verbalize their reasons in forms that can be converted to rules; that means they cannot rely on common sense for an answer; or if they do, then someone must be able to analyze their cognitive processes into rules.

Medical expertise clearly fits these constraints. Political savvy probably does not. And many areas of expertise fall in between.

CADUCEUS has a very broad expertise. It covers hundreds of diseases in internal medicine. It takes into account thousands of symptoms and lab tests. It does professionally well against diagnostic tests of its abilities. The home computer will truly be here when this fine program is available for everyday use. However, this represents the result of 10 years full-time commitment by an expert internist and the help of many colleagues and students.

MYCIN is much narrower, but deeper. It deals only with the diagnosis and treatment of organisms that infect blood or cause meningitis. In direct tests, it has achieved performance comparable with experts in the field (Buchanan, 1981). Furthermore, it has spawned a whole family of expert systems.

By taking away the specific medical knowledge domain, a shell of essential or empty MYCIN has been formed, called EMYCIN. Given a new problem domain, EMYCIN provides help in constructing an expert system to solve the problems. EMYCIN assumes an appropriate representation format for the new knowledge base (production system rules) and goal-directed backward chaining as

the appropriate inference mechanism. Production systems are simply collections of IF-THEN rules. Goal-directed backward chaining is the same process you use when you decide to get 40 dollars out of the bank because you are going to dinner. Other forms of knowledge representation and other inference mechanisms are available, but these have so far proved to be most useful.

Combining EMYCIN with new knowledge domains has yielded useful expert systems for pulmonary diseases (PUFF), blood clot diseases (CLOT), structural analysis (SACON), filling in and debugging these expert systems (TEIRESIAS), and explaining and tutoring these systems (GUIDON).

GUIDON (Clancey, 1981) is of special interest to us as this expert system is for instruction and training. In a project funded by ARI, Clancey at Stanford's Heuristic Programming Project has converted MYCIN into NEOMYCIN. NEOMYCIN reconfigured the rules in MYCIN so that they became more psychologically plausible as the basis for reasoning diagnostically. This allows GUIDON to provide test cases to students, allow them to analyze NEOMYCIN's diagnoses, and by interpreting the students' analyses, build up student models of their knowledge base and diagnostic reasoning abilities. Obviously, GUIDON represents a very large step forward in individualizing instruction and getting the computer to become a tutor.

A discussion of intelligent computer-aided instruction (ICAI) and tutorial systems would be utterly inadequate without mention of the dramatic and pioneering work of John Brown and his colleagues at Xerox PARC. In a now lengthy series of impressive tours de forces, Brown (1983) and company have provided imaginative demonstrations of how to change our mind set about computers in instruction. His fundamental thesis is that these sophisticated symbolic environments provided in personal LISP stations or Smalltalk environments can expand the use and effectiveness of learning-by-doing through discovery and apprenticeship.

Learning-by-doing environments can provide a tremendous opportunity for students to learn about themselves and their own thinking processes (metacognition) and their learning strategies. As an example, Brown and Burton capitalized on the current computer games mania by creating a computer coach for PLATO's "How the West Was Won" called simply WEST (Brown & Burton, 1982). WEST can watch two people playing each other and decide on its own when to interrupt politely. It can then point out to a player that, if he had only paid more attention to his opponent's moves, he might have discovered a strategy he did not possess. The obvious strength of this coach is that it can make use of the skills of the opponents without having to create unique scenarios for tutoring. It can in fact facilitate peer tutoring. ARI is sponsoring work to continue some aspects of this research and development.

Finally, I want to describe briefly a project called STEAMER (for a more detailed view, see Jim Hollan's contribution to this workshop). STEAMER has tried to integrate some of this research and development to provide intelligent computer-based instruction in propulsion engineering. This work arose directly from SOPHIE, another Brown, Burton, and de Kleer project (1982) providing explanatory qualitative models for training electronic troubleshooting for a relatively simple amplifier circuit.

309

In the first phase of STEAMER (Orlansky, 1983), a learning-by-doing environment was created with a sophisticated graphics interface that simulated thousands of components in a naval steam propulsion plant. This graphics package is a major product for instruction by itself. It may now be used within this Lisp environment for other kinds of instruction, or other LISP environments, such as LOOPS, may be converted to perform similar functions. This part of STEAMER provides a unique audio-visual instructional aid, and it is being used just for that purpose in a Navy setting.

The second phase of STEAMER is now underway to build a computer coaching facility into the same environment. Two strikingly different approaches are being used. One carries out a sequential dependency analysis of appropriate steps in any procedure of running the plant, i.e., this meter must be checked before that valve is opened. This leads to the sort of tutorial feedback that says, "Sorry, you should have checked A before doing B; you just blew up the ship."

The second form of tutorial tries to avoid the complexities of understanding the functioning of each single component and tries to create a deeper conceptual understanding of the causal connections among things. For instance, it may provide a moving graphic of a steam valve that can be controlled with all its functioning parts in display. A student can manipulate this model in gross ways, like increasing the pressure at a point and see, visually and dynamically, what the consequences are. Hopefully, this kind of learning, while it does not provide engineering competence, does provide the kinds of mental models that allow expert technicians to operate and troubleshoot the system effectively.

None of these ICAI projects have been evaluated in any rigorous fashion. In a sense, they have all been toy systems for research and demonstration. They have all raised a good deal of excitement and enthusiasm about their likelihood of being effective instructional environments. With the dramatic decreases in computer hardware costs we have recently seen, delivery vehicles for systems like these in the $10,000 range are a reasonable estimate. Clearly, it is the judgment of many that implementations of these systems in a variety of environments will prove to be highly cost effective. The real difficulty in making application decisions is in foreseeing the role that they will play as stepping stones to the future: to imbedded training devices; on-the-job computer aids for maintenance; personnel assessment devices that simulate real tasks; and to devices that actually carry out real tasks, such as imbedded maintenance or robotic control. For even if these systems are not yet cost effective, the only cost effective way of eventually deploying these systems may be to start using primitive and even unwieldy systems now. This is a difficult but very real problem.

## Cognitive Science Technology

The productive application of artificial intelligence to training and instruction rests heavily on the technology base of understanding provided by good cognitive science research. This research is advancing rapidly to provide a richer theory of learning that can actually be used in ICAI environments. Two provocative components of this research provide some insight into how ICAI can

be used to provide superior instruction and training. The first component to be described will be an analysis of misconceptions, and the second will be the use of learning strategies.

## Misconceptions

Consider the path a projectile will take after emerging from a curved tube at high speed. For both shapes in Figure 1 the correct answer is a straight line, tangent to the curve. Notice that more people said the exit path would be curved for the spiral tube than the semicircular tube. It is as if the extra turns were more likely to impart their curved "momentum" or "force" to the moving object. In fact, this bears some similarity to medieval impetus theory and Aristotelian ideas about motion. Other evidence strengthens this suggestion. For instance, similar misconceptions can be found for dropping things from airplanes, tossing coins, and other problems with moving objects. The striking fact about this finding is that all these people are high school graduates who have taken physics.

The second area of misconceptions deals with equations (Figure 1). Any relationship that can be represented algebraically can be expressed in words. Clement, Lochhead and Soloway (1980) report that 37 percent of a group of engineering students wrote 6S=P for "There are six times as many students as professors." Apparently, one source of this problem is that students deal with the equal sign as a static rather than a functional symbol. It is easy to imagine a room with 6 Ss and 1P and use the faulty equation to describe it. It is much more difficult to understand that this description calls for an operation--multiplication of P by 6--to achieve the goal of equalizing the P side of the balance.

Both misconceptions proclaim that simple errors may arise from complex and detailed systems of thinking that may in fact be transitional to more complete ways of understanding a knowledge domain. That is to say, the errors may not just be random, but arise systematically from erroneous ideas or misconceptions that people hold about an area of knowledge. Just such a view of errors has been studied within arithmetic by Brown and Burton (1982) and used as a diagnostic aid for prescribing remedial instruction.

The analysis of such misconceptions is proceeding in other knowledge domains. Identifying a misconception is a very difficult, time-consuming task; but it is well worth the effort because instruction can then be explicitly designed to deal with it and produce longer lasting and more immediately effective instruction. From the perspective of the present talk, it also has the powerful advantage that this kind of analysis is just right for ICAI systems that have to be explicit about their diagnostics and forms of remediation.

## Learning Strategies

Cognitive science is contributing important knowledge about how formerly mysterious components of thinking, like intuition, inference, and problem solving, are carried out. More important, there is every indication that
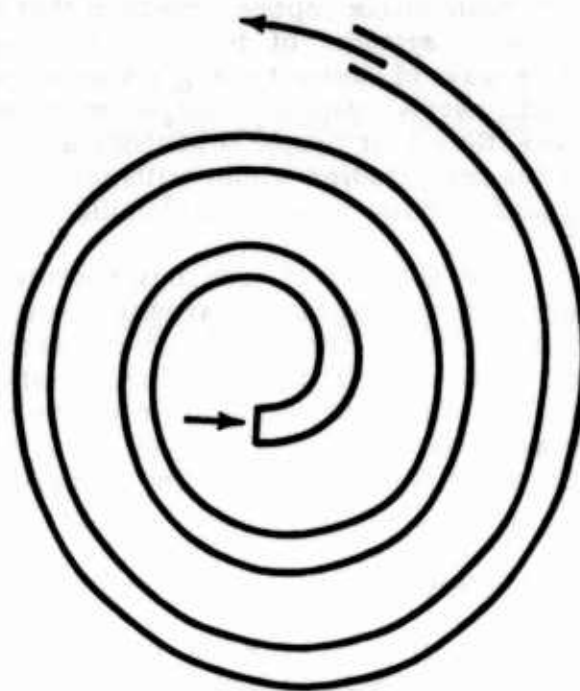
311

# MISCONCEPTIONS
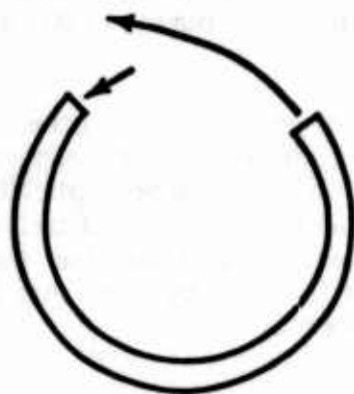
WRITE AN EQUATION FOR:
"THERE ARE SIX TIMES AS
MANY SOLDIERS THAN
CIVILIANS IN THIS ROOM."

$$6S = C$$

30%

60%

Medieval IMPETUS Theory

Figure 1.

these mental processes can be taught explicitly, although not easily (Nickerson, 1983; O'Neil, 1977). It may take the special audio-visual and explanatory powers of a sophisticated computer system to realize the full potential of the learning and teaching strategies that are being uncovered.

A short paper like this can only begin to give a flavor for the vitality of the research going on in this area, and its importance for ICAI. Much of the important work deals with working out the details between experts and novices, in terms of the strategies and heuristics they use, the structure of their knowledge, and guiding metaphors and analogies of their thought. By understanding these relations it is hoped that procedures for speeding up the expert-novice transition can also be found.

Experts carry out many of their activities automatically, below their general level of awareness, with little cognitive overhead or interference with other tasks. Speeding up the automaticity of these procedures in novices is something that can be done by computer after an effective cognitive task analysis. ICAI can help decompose tasks into manageable components so that not only are the tasks easier to learn, but the general skill of breaking things into bite-sized chunks can also be learned and used in many novel situations.

Most intellectual processes are not subject to awareness or introspection. This has the unfortunate effect that most of us are not well aware of how good our memory is; how many rehearsals it takes to remember a telephone number; what our misconceptions about physics, public policy, or people are; how well we can use imagery, analogies and metaphors; and a host of other mental skills. The study of metacognition is exploring how to make these facts of our mental life visible to us, so that we can begin to change and improve them.

Visual representations (graphs, charts, tables, working sketches, even cartoons) are all relatively underused when people try to learn and understand new material. Computers, with their dramatic display capabilities provide an environment that opens up rich new possibilities for using these tools to enhance our understanding and create new ways to learn.

Finally, learning without computers is often all too passive. Intelligent CAI can help ask the questions that test the boundaries of our understanding; provide counter examples to explode the limitations of overly narrow comprehension; make new connections among the things we know; and provide a flexible and generally supportive but challenging environment in which learning can take place.

## Conclusion

To do all this, the computer would indeed become a tutor, and replace to some extent those who would now be trainers and teachers. However, this development is still a long way off. Until then, however, ICAI systems will be able to take on an increasingly heavier burden, becoming more intelligent as the growth in technological power, cognitive science, and artificial intelligence

applications permits. This may be a Pollyannic vision of the future, but it is certainly one that must be taken into account in analyzing current directions and policies for training technology.

NOTE: The views expressed in this paper are the author's and do not necessarily reflect RI's official policy. Arpanet address: Psotka @ NPRDC. Mail address: Attn: PERI-IC.

## References

Brown, J. S., & Burton, R. R. An investigation of computer coaching for informal learning activities. In Sleeman, D. and Brown, J. S. (Eds.) Intelligent tutoring systems. London, England: Academic Press, 1982.

Brown, J. S., Burton, R. R., & de Kleer, J. Pedagogical, natural language, and knowledge engineering techniques in SOPHIE I, II, & III. In Sleeman, D., & Brown, J. S. (Eds.), Intelligent tutoring systems. London, England: Academic Press, 1982.

Brown, J. S. Learning by doing revisited for electronic learning environments. In White, M. A. (Ed.), The future of electronic learning. Hillsdale, NJ: Lawrence Erlbaum, 1983.

Buchanan, B. G. Research on expert systems (Report # CS-81-837). Stanford, CA: Dept. of Computer Science, Stanford University, 1981.

Clancey, W. J. The epistemology of rule based expert system: A framework for explanation (Report # CS-81-896) Stanford, CA: Dept. of Computer Science, Stanford University, 1981.

Clement, J., Lochhead, J., & Soloway, E. Positive effects of computer programming on students' understanding of variables and equations. Proceedings of the National Association for Computing Machinery Conference, Nashville, October 1980.

Feigenbaum, E. A., & McCorduck, P. The fifth generation. Reading, Massachusetts: Addison-Wesley, 1983.

Nickerson, R. S. Computer programming as a vehicle for teaching thinking skills. The Journal of Philosophy for Children, 1983, 4, 42-48.

O'Neil, H., Jr. Learning strategies. New York: Academic Press, 1977.

Orlansky, J. Have we increased the productivity of training? Proceedings of the Society for Applied Learning Technology, July 13-14, 1983.

Stevens, A., & Steinberg, C. Project STEAMER (NPRDC Tech. Note 81-21). San Diego: Navy Personnel Research and Development Center, 1981.

## ABOUT THE AUTHOR

Dr. Joseph Psotka is team leader of the Army Research Institute's Computer-Based Systems Development Team. In this position, he is involved with the application of artificial intelligence to answer the Army's needs in combat, doctrine, training and instruction. Previously, he was Senior Associate at the National Institute of Education and taught at the University of Waterloo. Dr. Psotka received his Ph.D. from Yale.