LETTERS

DARK AGES PANEL CORRECTIONS

Editor:

I was taken aback to note that the *The Dark Ages of AI* panel discussion at AAAI-84, in which I was one of the panelists, appeared in the Fall, 1985 issue of *AI Magazine*, since due to some communication gap I wasn't aware that the panel discussion was going to be published and I hadn't had a chance to proofread my section of the transcript. I was rather unhappy when I read the section that contained my remarks: Perhaps because of an accent that would not vanish after 20 years in this country, my remarks were, in significant places, embarrassingly garbled by the transcriber ("the most performed paradigmatic change"?, "AI has been the whole expectation of the problem"?, "Knowledge use invalidities has been the cause of misunderstanding"?), and in other places, the crucial "not" had been omitted or added, completely changing my intended meaning, "not" being generally very unforgiving in this regard (where I had said, "The problem is underestimation of the problems of multiplicity of generic knowledge structures," "is" appears as "isn't;" I am pretty sure I didn't say, "I also believe that faster architectures could do the trick," since at that stage in my talk, I was *criticizing* the belief that what it takes is faster architectures, while crucial epistemic problems remained unsolved). Perhaps it is best to outline the main points of my panel presentation to make clear what I really said (this time without an accent and slowly):

1. AI has already made significant paradigmatic contributions by fostering the idea of cognition as computation. This notion is bound to have far-reaching consequences to philosophy and psychology.
2. AI has already created a *style of programming*, viz., the family of frame-based systems with embedded procedures. This is a weak theory of mind (or mental architecture) in the sense that it says something about organization, but doesn't make any strong commitment about content. This is a *new* style of thinking and organizing computation and definitely very productive.
3. There has been a discontinuity in AI work that has caused the spurt in interest, and that is the emphasis on the role of *knowledge*. But because of the very weak theories of knowledge that we have, and an inability to characterize what kinds of problems can be successfully handled by these theories of knowledge, excessive expectations have been created. There is a poor understanding of the *multiplicity* of knowledge structures and processes that are operational in human intelligence. Instead of focusing on this, there is instead a tendency to believe that all it takes is faster architectures.
4. There *is* a body of commercial applications that can be thought of as "knowledge-rich" and "problem-solving-poor," where merely organizing the knowledge and making intelligent access possible can be very useful and profitable. The uniquely AI-style of programming that I alluded to above will contribute heavily towards these types of systems.
5. In general, good applications of AI require strong epistemic analyses, the skill for which is difficult to acquire and to teach.
6. All this interest in AI has certainly brought to light a number of hard and interesting research issues—which would not have come about but for the forced concentration on real problems in real domains. Projects such as the DARPA Strategic Computing Initiative will result in significant research and application advances, even though they may not be exactly in the forms which motivated the establishment of the initiatives.
7. AI is in a good situation where good research can be done in the context of building real systems for real problems. In fact, the real world problems give the constraints that otherwise are absent from purely "theoretical" endeavors.

B. Chandrasekaran
Laboratory for AI Research
Department of Computer
and Information Science
Ohio State University
Columbus, Ohio 43210

RUDE vs. COURTEOUS

Editor:

In vol. 6, no. 3 of the *AI Magazine,* the one known as the JUMBO edition (for either of two equally obvious reasons, neither of which is acronymic), Jack Mostow responded to my brief sketch of a performance-based paradigm for AI system development—the RUDE methodology. Whilst not aspiring to become the major contribu-

tor to some future collection of letters entitled, "The History of Hacking: a perspective on AI system development," I would like the opportunity to point out very briefly both a misconception and a major problem with Mostow's response.

First, he states that "Modifying an implementation directly (*i.e.,* RUDEly) is like patching compiled code." That is of course not a good thing. It is also not true, at worst it is patching compilable (or interpretable) code, *i.e.,* a machine executable specification. There are two problems here: the level of available machine executable specifications—current programming languages are cluttered with implementation details; and the strategy of modification—it should be more structured than "patching" suggests. If the dream of automatically replaying the derivation from specification to implementation becomes a reality then Mostow's COURTEOUS scheme will also be based on the modification (hopefully not patching) of compilable code. So perhaps it is not such a bad basis given less-cluttered machine-executable notations and a thoughtful modification strategy (elsewhere I have advocated a scheme of 'controlled' modification). Both improvements could contribute to a disciplined RUDE-based methodology for AI system development; I stated the dire need for such development in my original letter.

So much for the misconception. The major problem derives from the observation that AI problems are largely defined, and have to be evaluated, in a performance mode. How can we avoid "decompiling" (a nasty problem as Mostow points out) if the driving force behind modification is observation and evaluation of a system's behavior? We 'see' the system's inadequacies in terms of the behavior of an implementation, and must derive the implications for the specification, *i.e.,* we must "decompile". In fact the problem is somewhat more complex because the terms 'specification' and 'implementation' are really just convenient fictions: they are not totally separable concepts, and any specification can usually be re-specified more abstractly or more concretely. The same is true for implementations. A complex problem is characterized by a series of specifications differing in level of abstraction. More importantly for current purposes, comprehension of complex problems is facilitated by such a sequence of specifications a fact exploited by the well-known software engineering technique of stepwise refinement. Thus effective modification of a machine-executable specification is not just a matter of decompiling to obtain the abstract specification. We need a whole series of such decompilations.

I favor a development of the RUDE methodology that involves a technique of stepwise abstraction (Partridge, Procs. ECAI, Pisa 1984) to facilitate the decompiling which I see as necessary for AI system development. Decompilation, or stepwise abstraction, is obviously a hard problem, but remembering the drunk who searched under the street light for the keys that he had dropped in a dark alley, the hard route may be the only way to a solution. For although to Mostow "it seems much easier to support the implementation process than to invert it," the COURTEOUS methodology he proposed actually involves both problems: the hard problem of decompilation, and the perhaps very hard problem of automating a significant portion of a generalized design and implementation scheme.

I am convinced that development of a disciplined methodology based on the RUDE cycle is the best general approach to the problems of AI system design, but I remain unconvinced that the COURTEOUS route is the particular one to take. AI is, unfortunately, a behavioral science.

Derek Partridge
Computing Research Laboratory
New Mexico State University
Las Cruces, New Mexico 88003

KNOWLEDGE AND POWER

Editor:

"In Europe everything matters but nothing goes. In America nothing matters but everything goes."

Weizenbaum's *Computer Power and Human Reasoning* is as sparse an exception to the motto as is the recent letter by Mr. Kornell (Vol 6, No. 2). This letter contains "mind splitting" questions that produce a fair amount of confusion. It is annoying to ponder these moral issues when one is neck deep in AI trying to fix an exponential leak.

The two main questions were: "To what extent and in what ways is knowledge power?" and "... what is our responsibility as a community of researchers regarding the development of military AI technology?"

My reflex reaction was "Why bother at all?" Mr. Kornell has blocked this answer by sketching a scenario where AI is not leading to the "democratization of technology" but instead sharpens and widens the distance between the haves and the have-nots—within a society as well as between the nations.

Applicability of his scenario within the American society is debatable. The democratic tradition might prove to be too strong for AI to exercise a fundamental influence on the structure of society. To argue otherwise may be interpreted as arrogance. In my opinion, AI applications are the next wave of "toys" rippling by; favoring a few, disfavoring others, rearranging the *status quo* for many, but not fundamentally affecting the distribution scheme of wealth and power.

In contrast, a variant of the scenario seems to be highly plausible on the international level. I don't know the exact figures, but I estimate that less than 5% of the world population consumes more than 50% of the yearly world pro-

duction. In addition, this imbalance widens, if not through "American shrewdness" (from the perspective of the international outsiders) then through population growth that surpasses economic progress. This process has already unfolded for decennia. The stakes are high for the U.S. AI is simply the next device employed to defend this skewed international wealth and income distribution.

Kornell's question about the responsibility of the AI community regarding the employment of military AI technology is therefore out of date. American *communis opinio* considers its wealth the fruit of hard, ongoing labor. Whether this wealth needs a defense is an esoteric question that may rage in academic circles but its conclusions are not likely to be heard outside the ivory towers.

Given this state of affairs I prefer to wonder whether we can influence the processes on the other side of the fence. This leads to questions like:

"Is there a way that AI can leverage undeveloped potentials in the Third World?"

"What is an inducement scheme for such work and what would be the proper mechanism for a transfer of results?"

Obviously, these questions can be modified by replacing AI's role with other segments of America's R&D community. Such questions must have been posed many times over. Answers provided have, by and large, as yet not worked out.

Is there reason to believe that AI might succeed where others have failed? Or is there a unique set of historic circumstances that prohibits a balanced global distribution of wealth?

Dennis de Champeaux
1242 Redcliff Drive
San Jose, California 95118

AWARD NOMINATION

Editor:

In the midst of all the comments you're probably receiving about how overcrowded IJCAI has become, I'd like to nominate Claudia Mazzetti, Linda Quarrie and Phyllis O'Neil for the Battered Pot Award. The winner is to receive a slightly battered pot, three liters of hot water and a box of Epsom Salt. The runners-up are to receive a supply of Excedrin sufficient to last to IJCAI-87. The award may be modified as needed according to the state of exhaustion of the recipient.

Steven C. Blake
Canoga Park, CA

For the record (and the relief of the nominees), the AAAI will not be co-sponsoring IJCAI again until 1989.

—Ed.

Hitech wins North American Computer Chess Championship.

Hitech, CMU's chess computer / program, has won the annual ACM sponsored North American Computer Chess Championship. In the tourney, held in Denver from October 13-15, Hitech defeated all its major rivals to achieve a perfect 4 0 score. Although the sudden advent of this new contender may seem surprising, it should be noted that at the time of the tournament Hitech was already the highest rated chess computer in the World. To that time, Hitech had played 21 tournament games against human competition and achieved a record of 16.5 4.5, giving it a US Chess Federation rating of 2233, the highest ever achieved by a computer. Prior to the ACM event Hitech won a tournament with four Masters in it, the first time such a feat had been accomplished.

In the ACM tourney, Hitech defeated all its major competitors. In round 1 it beat LACHEX, a program from Los Alamos labs running on a Cray-1. In round 2, it beat PHOENIX from the University of Alberta running on several multiplexed VAX's. In round 3 it defeated the eventual runner-up, Bebe, a bit-slice special purpose machine from Sys-10 in Chicago. In round four it beat defending champion Cray Blitz running on a Cray X-MP 48.

The final ACM standings were:

Hitech	4 –0	Chaos	2 –2
Bebe	3 – 1	Lachex	$1\frac{1}{2}$ –$2\frac{1}{2}$
Intel SoftW	$2\frac{1}{2}$ –$1\frac{1}{2}$	Spock	$1\frac{1}{2}$ –$2\frac{1}{2}$
Cray Blitz	2 –2	Ostrich	1 –3
Phoenix	2 –2	Awit	$\frac{1}{2}$ –$3\frac{1}{2}$

This was the first time in 4 years that Cray Blitz had lost to a computer, and it lost to both the first and second place winners.

CMU's Hitech is based mainly on a very fast special purpose architecture that involves a move generator made of 64 specially designed VLSI chips. This "Searcher" is able to process 200,000 positions/second including evaluating the positions almost instantaneously. Evaluation uses specially designed hardware and compiled knowledge prepared by the software "Oracle" of the system, and down-loaded prior to the start of the search.

Hitech has been in preparation for nearly 3 years, during which time graduate student Carl Ebeling designed the hardware and MOSIS manufactured the special purpose VLSI chips. Since the beginning of 1985, a team headed by Hans Berliner has fleshed out the system to be a useful chess playing entity. Involved in the system building were: Carl Ebeling: Hardware design and construction, Gordon Goetsch: system software, Andy Palay: initial concept and search strategies, Murray Campbell: openings and testing; Larry Slomer: hardware construction, and Hans Berliner: chess knowledge.

Since its debut on May 25, 1985, Hitech has been progressing steadily up the USCF rating ladder at the rate of about 8 points/game. Its present rating is 2555, about 50 points higher than any computer has ever achieved.

Because it is possible to increase the knowledge of Hitech without slowing it down, we see no reason at this time for Hitech's progress to be slowed in the immediate future.

— Hans Berliner