M A S S A C H U S E T T S   I N S T I T U T E   O F   T E C H N O L O G Y

**ARTIFICIAL
INTELLIGENCE
LABORATORY**

9 March 1994

Ramesh Patil, Editor
AI Magazine
AAAI
445 Burgess Drive
Menlo Park, CA 94025

To the Editor:

On the topic of *Intelligence and Reason*:

In the dozen or so years of our acquaintance, I have invariably found Oren Etzioni's perspective to be insightful and enlightening. It was therefore with some dismay that I read his recent article ("Intelligence without Robots", *AIMag* Winter 1994), in which his otherwise excellent arguments in favor of software agent research are buried under a muddy (and unwarranted!) attack on robotics-based research. It is a shame that, in his response to Brooks's papers, Etzioni has resorted to precisely the sort of polemic that he seeks to discredit.

Etzioni has some good points to make. Software agents *are* useful and interesting. The challenges posed by the increasingly complex computational world necessitate a response from the AI community. Software agents appear to be a promising approach for coping with this potentially overwhelming complexity. And software agents do not rely on robotics research for their basic technology. From this, Etzioni reasonably concludes that robotics alone is insufficient as a basis for all aspects of AI research. Etzioni at his best sums up his own lucid arguments when he says that "softbotics and robotics are complementary methodologies for investigating intelligent agents in real-world environments."

Etzioni goes on to claim some of Brooks's arguments for his own. Brooks argues for architectural completeness, citing this as a reason for dismissing traditional decompositions of intelligence. But there can be little doubt that software agents *are* complete embodiments, within the computational world that they inhabit. Further, there are reasons to believe that some of the traditional AI techniques which Brooks decries but Etzioni insists "still have much to contribute to AI" may be particularly appropriate in the domain of software agents. Etzioni saves traditional AI.

Room 811, 545 Technology Square, Cambridge, Massachusetts 02139 • (617) 253–2663

For software agents, yes. But has Etzioni really shed any light on the thorny issue of cognitive architectures for creatures that inhabit the physical world? This is where Etzioni's rational and reasonable arguments wear thin and his claims become increasingly speculative or even, in the end, entirely ungrounded. As an example, consider Etzioni's "strong version" of his argument. Premise: One's research is shaped by in part one's test bed. Unarguable. Premise: Robotic test beds emphasize sensing, acting, and the hardware and software systems required to support them. Absolutely. Conclusion: Core AI issues are better studied in software agents. Where did this come from? Presumably software agents require detailed knowledge of software systems, network resources, and other arcane details that Etzioni may find familiar but the average electrical or mechanical engineer would consider truly bizarre. The conclusion, *even* if true, certainly doesn't follow from the premises. Architectural completeness—for agents implemented in software as well as in hardware—comes at a cost. Etzioni finds the costs more acceptable for software agents. Reasonable enough; let him conduct his research there—but without pretending that he is somehow getting away without overhead.

And then there is the question of what his research will tell us. Etzioni argues next that software agents admit investigation of the full spectrum of real-world problems. And who can argue with him when he insists that software environments are "real"? Don't we all use them each and every day? Aren't they where many of us, in some sense, "live"? Compelling, yes, and probably true when it comes to many of the "architectural completeness" claims that Etzioni makes for his software agents. But let us take a step back. After all, is there any environment *more* engineered, *more* artificial, than the world of software? The very concept of cyberspace is intriguing in large part *because* the information that is readily available and the actions that are easily undertaken are so unlike those in the "real" world. The fact that a software agent lives and breathes in this artificial world may make it a wonderful vehicle for exploring a whole host of problems, but it says nothing whatsoever about the ability of that architecture to function in *our* world.

Etzioni ends his "Engineering Methodology" section with the bold statement that "softbots satisfy every facet of Brooks's engineering methodology," but omits the caveat that they do so *only vis a vis the particular world in which they operate*. All agents are architecturally complete for the environments in which they reliably function. Brooks's criticism of previous software *and robotic* systems derives from the artificially engineered nature of the environments in which they functioned *coupled with the claim that these systems allow one to draw conclusions about the functioning of systems in the "real" world*. No one disputes the adequacy of software agents for a software world. But what do they tell us about *our* world?

So the first half of the paper ends having "showed that Brooks's methodological arguments actually support the softbot paradigm," more or less.[1] Further, as Etzioni now argues—and few would dispute—software agents are a useful and interesting research area of their own right. So why does Etzioni feel compelled to take issue with his *own* statement that "developing mobile robots as a basis for softbots is about as plausible as developing softbots as a basis for mobile robots"?

Perhaps the problem arises in his conflating of software issues ("for example, responding to error messages; cloning softbots on remote machines...") with cognitive architecture. Or maybe it is just that he finds the "real" world—with its problems ranging from sensor noise and occlusion to reasoning about shapes and physical actions—to be "a distraction". But somewhere along the line, Etzioni leaps from the reasonable premise that software agents are interesting and useful to the outrageous conclusion that robotics research is utterly unnecessary.

I am willing to follow Etzioni through his argument that software agents are an exciting research area. I agree with his suggestion that "many AI issues can be studied profitably in real-world software environments." But Etzioni fails to provide compelling reasons why we should expect lessons learned in the decidedly artificial world of the computer to carry over to the extremely different world in which we live. Perhaps they will. Perhaps, when the almost-but-not-quite-irrelevant field of robotics solves the input and output problems, Etzioni's software agent architectures will provide the cognitive components necessary to create human-like embodied intelligences. But there remains ample room for doubt.

Brooks espouses robotics research because he believes that the details of sensation and action are *not* peripheral to AI. What for Etzioni is irrelevant and distracting is for Brooks at the very heart of the issue. If Brooks is correct, software agent research may provide interesting and enlightening results, but it will never supplant the central position occupied by robotics in our efforts to understand human activity in our own very real world. But even if he is wrong, there is no reason to suppose that research into

[1]The logical argument in the final part of this section is a red herring, though the conclusions Etzioni ultimately draws about grounded symbols in the software world are reasonable. Certainly software agents's symbols are grounded *in their world*, though whether that world is itself grounded is both unclear and irrelevant. The diatribe on punctuated equilibrium as a theory of evolution seems equally unnecessary, particularly given the peripheral nature of the issue to Brooks's or Etzioni's claims. To be perfectly fair to Brooks, however, Etzioni might have done him the favor of analyzing (and presumably attacking) both the ethological and the developmental evidence as well. In any case, this passage affords Etzioni with the opportunity to make one of his most tellingly ironic statements: "We should not underestimate the amount of evolutionary change underlying our higher cognitive functions." True, but neither—given the inconclusiveness of the evidence—should we overestimate it.

software agents will not be enhanced by what Etzioni himself calls the "complementary methodologies" of robotics research. It would be nice to believe that the field of artificial intelligence is not utterly incompatible with reason.

Sincerely,

Lynn Andrea Stein
Class of 1957 Assistant Professor
Electrical Engineering and Computer Science