AI NEEDS MORE EMPHASIS ON BASIC RESEARCH

John McCarthy
Computer Science Department
Stanford University
Stanford, CA 94305
`jmc@cs.stanford.edu`
`http://www-formal.stanford.edu/jmc/`

1983

Abstract

Too few people are doing basic research in AI relative to the number working on applications. The ratio basic/applied is less in AI than in the older sciences and than in computer science generally. This is unfortunate, because reaching human level artificial intelligence will require fundamental conceptual advances. Even the applied goals proposed by various groups in the U.S., Europe and Japan for the next ten years are not just engineering extrapolations from the present state of science. Their realization will require more basic research than is now being done.

Jon Doyle put it this way in a recent net message. "... tentative, but disturbing conclusion: that the students interested in AI are not very interested in fundamental questions, open problems, and long term research, but instead are eager to get in on big, build-it-now projects in expert systems and natural language interfaces." He was definite about CMU, but he conjectured that the situation was similar elsewhere, and I suppose student preferences are similar in different places.

I'll begin with a few recriminations and then try to be more constructive. First the Government, specifically DARPA and NSF, had a fit of extreme ‘practicality’ in the early 1970s. The Mansfield amendment required DARPA to claim short term military relevance for what it supported, and NSF diverted much of its resources to ‘Research Applied to National Needs’. The older sciences were able to resist this in NSF but lost their DARPA support completely. AI, which was more dependent on DARPA than the others were, survived but wounded. The situation has improved in both places in recent years.

Second the opportunities to make money have perhaps lured some people away from research per se. I don’t really know the extent to which this is true. Maybe they were tired of research.

Third much of the theoretical work in AI is beside the point and unlikely to lead to advances toward human level intelligence. The mathematically talented like well-defined conjectures the wherein the mere statement of the result that has been proved or the asymptotic behavior of the algorithm discovered wins instant scientific recognition. AI badly needs mathematical and logical theory, but the theory required involves conceptual innovations - not just mathematics. We won’t reach human level intelligence by more algorithms reducing the complexity of a problem from n^2 to $n \log n$ and still less by proofs that yet another problem is unsolvable or NP-complete. Of course, these results are often very significant as mathematics or computer science.

Fourth, like many fields AI is given to misguided enthusiasms in which large numbers of people make the same errors. For example, much of the present work in natural language processing seems misguided to me. There is too much emphasis on syntax and not enough on the semantics. Natural language front ends on programs that convert between existing AI formalisms and English miss the point. What we can learn from natural language is not how to express in English what we already know how to express in computerese. Rather we must study those ideas expressible in natural language that no-one knows how to represent at all in a computer.

We also won’t reach human level intelligence by building larger and larger production systems involving more and more facts all on the same level. Of course, these systems of limited intelligence may have substantial practical utility.

Now that I’ve finished grumbling, I’ll try to be constructive.

1. People beginning their research careers should think about the long term goals of AI and should think how to apply their own talents in the best way. If they can do first class basic research they should.

2. In my opinion, the key problem at present is the formalization of common sense knowledge and reasoning ability. It still looks to me that separating epistemology from heuristics will pay off.

3. We need to think hard about how to make experiments that are really informative. At present the failures are more important than the successes, because they often tell us that the intellectual mechanisms we imagined would intelligently solve certain problems are inadequate.

4. We need good problem domains — the AI analog of what the Drosophila did for genetics. The Soviet computer scientist A. S. Kronrod once referred to chess as the *Drosophila* of artificial intelligence, because it permitted comparison of human and artificial intellectual mechanisms. Unfortunately, chess was discouraged as a serious problem domain, and most chess programming is carried on at the level of sport rather than science. In particular, there is little publication about the intellectual mechanisms involved, and the race often involves merely faster hardware.

5. I also believe there is a large payoff in a more general analysis of the concept of pattern.

Finally, let me cheerfully admit that general harangues like this one are no substitute for scientific papers setting forth specific problems in detail.

I hope that other members of AAAI will express their own opinions about what the basic research problems are.

/@steam.stanford.edu:/u/ftp/jmc/basicresearch.tex: begun 1996 May 14, latexed 1996 May 14 at 8:47 a.m.