

Reconsiderations

Nils J. Nilsson



Nils Nilsson.

■ In 1983, I gave the AAAI president's address titled "Artificial Intelligence Prepares for 2001." An article, based on that talk, was published soon after in *AI Magazine*. In this article, I retract or modify some of the points made in that piece and reaffirm others. Specifically, I now acknowledge the many important facets of AI research beyond high-level reasoning but maintain my view about the importance of integrated AI systems, such as mobile robots.

In 1983, I gave the AAAI president's address titled "Artificial Intelligence Prepares for 2001." An article, based on the talk, was published soon after in *AI Magazine*.¹ Here, with the benefit of years, I retract or modify some of the points made in the article and reaffirm others.

A prepublication draft of the article, circulated among my colleagues in the AI Center at SRI, evoked a spirited rebuttal of those parts that limited the scope of AI and that described the role of logic and logical inference in AI. I encouraged the authors of the rebuttal, Sandy Pentland and Marty Fischler, to submit it for publication in the same issue of *AI Magazine*.² They did, and I belatedly acknowledge that many of their criticisms were well taken.

In the article, I claimed:

... not *all* symbolic information processing contributing to intelligent, perceptive behavior is of a piece. There are joints at which to carve this large body of science and engineering, and there are other disciplines that have already established legitimate claims on some of the pieces.

Imagining that these "other disciplines" would pursue their "claims," I proposed to limit AI to:

... what might be called *high-level* reasoning and perception. This core is mainly concerned with the collection, representation, and use of *propositional* or *declarative* knowledge. (Such knowledge is of the type that can be stated in

sentences of some form, as contrasted, say, with knowledge that is implicit only in procedures or in ad hoc data structures.)

Pentland and Fischler countered that:

AI research has *already* defined for itself a set of “core topics”: the study of the computational problems posed by the interrelated natural phenomena of reasoning, perception, language and learning. These phenomena may, of course, be viewed from many other vantage points including those of physics, physiology, psychology, mathematics and computer science. AI has continued to survive as separate from these other sciences because none of these other disciplines focus on developing computational theories for accomplishing intelligent behavior.... Their central interests remain quite different.

Quite true! It was imprudent to try to limit AI to just the high-level-reasoning part of intelligent behavior. AI has quite properly taken on the whole job, which is just as well because none of the other disciplines is doing it. “Human-level AI” (which I discuss elsewhere in this issue) will require “reasoning, perception, language, and learning” and more. In fact, for the past fifteen years or so I have been investigating architectures³ and formalisms quite distinct from high-level reasoning.

Even so, I still think that declarative representations (and manipulations of them) constitute an extremely important part of AI. There is a need for intelligent agents to have and to be able to use declarative knowledge in addition to the “procedural knowledge” that is encoded in special-purpose routines. Only those programs and circuits in which procedural knowledge is embedded can employ it. Declaratively represented knowledge, on the other hand, can be used for a wide variety of more general purposes—some of which might be unforeseen when the knowledge is installed or otherwise obtained. Also, in order to interact with humans at a high level of understanding, agents will need to be able to respond appropriately to declarative statements and commands—such as “Rooms on the second floor can accept deliveries only on Tuesdays.” Additionally, we want agents to be able to learn from books and other declaratively expressed material.

I claimed in the article that declaratively represented knowledge should be represented “in some sort of propositional, logic-like, formalism—and that much of the manipulation of that knowledge can and should be performed by mechanisms based on logical operations and rules of inference.” Pentland and Fischler quite appropriately objected to this version of what I called “the propositional doctrine.” Again with the benefit of years, I would now include many representational forms under the

heading “declarative knowledge.” Some examples are logical sentences, Bayes networks, episodic memory, and case studies. I would also include under the heading “reasoning,” many kinds of computation over these forms (in addition to the standard logical rules of inference).

And nowadays, I wouldn’t necessarily place reasoning at the “top” of an AI architecture—communicating with “lower processes” through procedural attachment. Perhaps the decision by an agent to reason should itself be governed by procedural mechanisms.

Pentland and Fischler complained that “... AI research often concentrates on the formal aspects of knowledge representations to the exclusion of how the representation’s symbols are embedded in, and derive meaning from, the world.” They say, “It is the embedding of a formal system in the world that gives meaning to the system.” I agree. Whatever is represented declaratively must be linked, ultimately through sensors and effectors, to whatever “world” the system behaves in. (Presumably, much of what is represented procedurally is already so linked.)

Let me turn now to some points made in the article that I take this opportunity to restate and reaffirm.

I stated that it was far too early to agree with the claim (often attributed to Marvin Minsky) that intelligence is a kludge—too messy to yield to “neat” theories about it. I still maintain what I said then, namely:

As scientists and engineers, we should continue to attempt to simplify, to organize, and to make elegant models—otherwise there are serious doubts that we would ever be able to understand enough about intelligence to design intelligent machines or to teach these design methods to students. If bridges had to be kludges, we wouldn’t have a man-made bridge across the Golden Gate because complex bridge-building couldn’t be understood, taught, or remembered. Successful engineering requires the frictionless case and a succession of gradually more complex models. I think AI has been reasonably successful so far in inventing understandable and useful theoretical frameworks and that it would be inappropriate for us to discontinue these attempts.

Related to the “intelligence-is-a-kludge” thesis is the tension between “neats” and “scruffies.” I still believe:

A dynamic field needs scruffies (informed but not overly inhibited by neat theories) at its expanding frontier of knowledge, and it needs neats to codify, clarify, and teach its core concepts. A field that is scruffy to the core has simply not yet matured as a science, and one that does not have a scruffy exterior is simply sterile.



Nilsson's Original Address Appeared in this Issue of AI Magazine.

His presidential address was delivered on 11 August 1983 at the National Conference on Artificial Intelligence in Washington, D.C.

Near the end of the article, I made a suggestion that I think still has much to recommend it:

A project should be initiated whose goal is to develop a class of new AI programs that would have a continuing existence. . . . let's call them *computer individuals*. The ground rule would be that they should never be turned off . . . They would have a constantly changing model of the world and of the user(s). They should be able to engage in extended dialogs in natural language. Designing such programs would stimulate (in fact force) research in machine learning because it would be manifestly unintelligent for a computer individual existing over time not to benefit from its experiences....

We can think of several varieties of computer individuals, including personal assistants, meeting schedulers, expert consultants, and mobile robots. . . . For the basic research purposes that such projects would serve, what specifically these robots would do is relatively unim-

portant. They could do anything that requires moving around in and sensing a real environment and manipulating that environment in some way. Aiming for roughly the same sort of sensing and manipulative abilities that people have would give us plenty of AI challenges.

In considering basic research projects of this kind, it is important to avoid the trap of insisting that the project be directed at some specific useful application. To attempt to justify robots by proposing them as useful highway constructors, for example, is misguided because general-purpose robots may be an inefficient way to solve the road-building problem—or the transportation problem—or any other specific problem. Any single application of robots alone is probably insufficient reason to justify their development. The whole reason for robots is their general-purposeness!

Finally, as a conclusion to this note, I repeat the introduction to the 1983 article:

Those of us engaged in artificial intelligence research have the historically unique privilege of asking and answering the most profound scientific and engineering questions that people have ever set for themselves—questions about the nature of those processes that separate us humans from the rest of the universe—namely intelligence, reason, perception, self-awareness, and language.

It is clear—to most of us in AI, at least—that our field, perhaps together with molecular genetics, will be society's predominant scientific endeavor for the rest of this century and well into the next . . .

Notes

1. Nilsson, N. J. 1983. Artificial Intelligence Prepares for 2001. *AI Magazine* 4(4): 7–14, Winter, 1983
2. Pentland, A. P., and Fischler, M. A. 1983. A More Rational View of Logic or Up Against the Wall, Logic Imperialists! *AI Magazine* 4(4):15–18.
3. Nilsson, N. J. Teleo-Reactive Programs and the Triple-Tower Architecture. *Electronic Transactions on Artificial Intelligence*, 5 (2001), Section B: 99–110. (Available from www.ep.liu.se/ej/etai/2001/006/.)
4. Nilsson, N. J., Teleo-Reactive Programs for Agent Control, *Journal of Artificial Intelligence Research*, 1: 139–158, January 1994. (Available from www.cs.washington.edu/research/jair/contents/v1.html/.)

Nils J. Nilsson is Kumagai Professor of Engineering, Emeritus of Stanford University. He is the author or coauthor of several AI textbooks and many of the seminal papers in the field. His career in AI began at SRI International in 1961 and continued at Stanford University where he was the chair of the Department of Computer Science from 1985 to 1990. Nilsson is a Fellow and a past president of AAAI. He received the IJCAI award for Research Excellence in 2003 and the AAAI Distinguished Service award in 2005. He can be reached at nilsson@cs.stanford.edu.