Epistemological and Heuristic Adequacy Revisited

MATTHEW L. GINSBERG

1269 University of Oregon, Eugene, OR 97403-1269
⟨*ginsberg@cirl.uoregon.edu*⟩

In what has become a landmark paper, McCarthy and Hayes [1969] observed that an effective AI program must deal with both epistemological and heuristic difficulties. The epistemological problems arise from the fact that the program must be adequate in *theory*: it must be able to solve problems given access to arbitrarily large computational resources. The heuristic difficulties stem from the fact that such resources never exist in practice; the program must, in fact, be able to solve problems given the computing power that is actually available. In terms somewhat more specific to AI, McCarthy and Hayes defined *epistemological adequacy* as the ability of a program to represent knowledge about the world and recognize valid chains of inferences drawn using this knowledge. *Heuristic adequacy* was defined as the ability actually to generate useful conclusions from input data. Continuing to refine these notions, McCarthy in 1977 wrote:

> The epistemological part of AI studies what kinds of facts about the world are available to an observer with given opportunities to observe, how these facts can be represented in the memory of a computer, and what rules permit legitimate conclusions to be drawn from these facts. It leaves aside the heuristic problems of how to search spaces of possibilities and how to match patterns. [p. 1038]

In both papers, McCarthy and Hayes go on to suggest that it is the epistemological problems that should be the focus of AI research. The argument, roughly speaking, is that epistemological issues are a separable subproblem whose solution will underlie subsequent work on heuristics and other practical techniques. This overall approach has driven much of the work on nonmonotonic reasoning[1] and other areas of knowledge representation since the mid-1970s.

The view I would like to defend in this short article is that McCarthy and Hayes are exactly wrong: epistemological and heuristic adequacy are *not* separable, and any attempt to solve either in isolation from the other is doomed to failure. This view is supported by an examination of both AI's recent successes and its relative failures.

Progress in Satisfiability and Constraint Satisfaction

One primary lesson of AI research over the past decade is that virtually every interesting AI problem is NP-complete or worse. This is clearly true of search pro-

[1] Nonmonotonic reasoning is the study of inference methods where the set of conclusions grows *nonmonotonically* with the set of assumptions. The standard example involves the fact that a bird Tweety will generally be concluded to fly. The assumption set gets *larger* is we add the fact that Tweety is an ostrich, but the set of conclusions presumably gets *smaller* because our belief that Tweety can fly will be retracted.

grams such as those that play games or prove theorems, and has also proven true of natural language and vision work—recognizing objects in a scene involves understanding the scene as a whole and what it is likely to contain, and processing a sentence involves both search *and* broad domain knowledge. A somewhat introspective example of Lenat's involves using context to distinguish the utterance, "It's very hard to wreck a nice beach" from, "It's very hard to recognize speech."

AI's fundamental dependence on NP-complete problems implies that the heuristic concerns involved in solving these problems are central to the AI enterprise as a whole. Furthermore, the subfield of AI that focuses on the solution of NP-complete problems (satisfiability and constraint satisfaction) is arguably the subfield that has seen the most progress in the last decade. There has been a host of new theoretical techniques introduced: Harvey's limited discrepancy search, Selman's GSAT family of algorithms, and my own dynamic backtracking are a few examples.

Experimental work has gone hand in hand with these theoretical advances. Indeed, algorithmic suggestions are no longer viewed as viable unless there is significant experimental evidence that they actually work. Where early AI experimental work involved tens or perhaps hundreds of problem instances with search spaces of size 2^{30}, modern work involves thousands to hundreds of thousands of problem instances of size 2^{100} to 2^{2000}. The observed performance of AI systems is improving rapidly, and the algorithmic work mentioned in the previous paragraph is almost entirely an outgrowth of insights gleaned from experimentation. It is a search for heuristic as opposed to epistemological adequacy that is driving progress.

The Relative Lack of Progress in Nonmonotonic Reasoning

Finally, I would like to consider one of AI's more disappointing research programs, the relative lack of progress in nonmonotonic reasoning work. Most of the work on nonmonotonicity over the past decade has focused on two specific toy problems: inheritance hierarchies and reasoning about action. Progress has been spotty, and many epistemological problems remain.

This theoretical work fails to address the original motivation for the introduction of nonmonotonic reasoning into AI. Nonmonotonic capabilities were supposed to make inference systems *faster* by allowing these systems to jump to reasonable conclusions quickly. In addition to the epistemological problems they have introduced, however, the computational properties of all proposed nonmonotonic reasoning schemes are far worse than those of their monotonic predecessors.

The problem seems to be that the space of possible nonmonotonic reasoning systems is immense. There surely exists some formalization that has attractive epistemological *and* heuristic properties, but we are not making much progress in finding it. A focus on (or even an interest in) the heuristic properties of new nonmonotonic theories would prune the search space immensely.

Summary

AI's conventional wisdom has been that epistemological problems should be solved first, to be followed by work on heuristically adequate implementations—the epistemological work informs the heuristic research.

In actuality, each element informs the other. AI's need for heuristic adequacy has strong epistemological consequences. The epistemological problems we are tackling are not greatly different from those with which philosophers have grappled for a millennium; our new perspective and insights are a consequence only of our inevitable focus on computation. The artificial intelligence community as a whole would benefit from recognizing this, and changing our common perspective to embrace the heuristic concerns that make our discipline unique.

ACKNOWLEDGMENTS

I would like to thank Tania Bedrax-Weiss, David Etherington, Moises Goldszmidt, and an anonymous reviewer for comments on a draft of this article; the views expressed here are mine alone. This work has been supported by ARPA/Rome Labs under contract F30602-93-C-00031.

REFERENCES

McCarthy, J. 1977. Epistemological problems of artificial intelligence. In *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, Cambridge, MA, 1038–1044.

McCarthy, J. and Hayes, P. J. 1969. Some philosophical problems from the standpoint of artificial intelligence. In *Machine Intelligence 4*, B. Meltzer and D. Mitchie, Eds. American Elsevier, New York, 463–502.