AI Magazine Volume 4 Number 4 (1983) (© AAAI)

A More Rational View of Logic or, Up Against The Wall, Logic Imperialists!

Alex P. Pentland
Martin A. Fischler
(The Loyal Opposition)
Artificial Intelligence Center
SRI International
Menlo Park, CA 94025

THE AAAI PRESIDENT'S ADDRESS (the previous article, by Nils Nilsson) presents an eloquent argument for a particular AI paradigm that may be summarized by what Nils calls the "propositional doctrine:"

> AI is the study of how to acquire and represent knowledge within a logic-like propositional formalism, and the study of how to manipulate this knowledge by use of logical operations and the rules of inference.

Although we concur with many of Nils's other assertions, this propositional doctrine seems far too extreme: a lot of interesting and important AI research is done outside of the logic-and-theorem-proving paradigm.[1] Indeed, the view that other lines of inquiry serve only to produce tools that may be procedurally attached to an AI (logic-and-theorem-proving) architecture seems a kind of Logic Imperialism to those of us they wish to relegate to working in the procedure factories.

This dismissal of other avenues of research as "not really AI" would normally be cause only for a knowing shake of the head and a small chuckle, but when such views are promulgated by the President of the AAAI it is time to take up arms against the logic-and-theorem-proving set—there is a danger that someone might actually take them seriously!

This paper, therefore, constitutes an initial salvo over (into?) their bow. We will focus on two central questions in this rebuttal:

- What is an appropriate research paradigm for AI?
- What role should logic-like formal languages and deduction play in the study of AI?

A Paradigm for Artificial Intelligence Research

Nils appears very concerned with establishing a unique niche for AI, and argues that the propositional representation of knowledge should be the core topic of AI—in part because no other discipline claims it. His concern seems, to us, both misguided and unnecessary.

It is unnecessary to set up a dichotomy between AI and other sciences by defining propositional formalisms as the 'core topic' of AI. 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

[1] A more formal characterization of the "logic-and-theorem-proving paradigm" discussed throughout this paper is summarized by the claim that (1) a formal logic language such as the predicate calculus and (2) formal tools such as resolution and unification — i e., syntactic operations that can be applied without consideration of a predicates' referent — are appropriate for the characterization and execution of general-purpose reasoning

as separate from these other sciences because none of these other disciplines focus on developing computational theories for accomplishing intelligent behavior. It should not bother us, therefore, if our study of intelligence borrows results, observations or techniques from these other disciplines; or if these other disciplines occasionally address some of the same problems, and use some of the same techniques. Their central interests remain quite different.

The propositional doctrine, however, is even *worse* than unnecessary. The central contention of the propositional doctrine is that intelligence should be studied by means of propositional, logic-like formalisms. An analogous statement would be "physics is the study of the laws of nature by use of partial differential equations (PDE's)." In the latter case it is clearly recognized that the goal of physics is to understand and formulate the laws of nature, and that PDE's provide a powerful and useful tool. It seems silly, however, to say that if you don't use PDE's then you aren't doing physics.

The *meaning* of the symbols being manipulated should be the prime concern, and not the syntax of the manipulation. Concentration on meaning rather than form is the crucial difference between a science such as physics, and a the formal branches of philosophy such as logic and mathematics. Science uses tools such as logic in order to establish a predictive isomorphism to the problem domain, while formal branches of philosophy concentrate on the properties of these tools. The preoccupation with logic-like formalisms evident in Nils's propositional doctrine has things exactly backwards: the *problem* should dictate the computational tools (as in science), rather than having the tools define the problem (as in logic).

The difference between AI as science and AI as a problem in logic is critical, because meaning (semantics) cannot reside within the formal systems of logic. Logic's inability to capture meaning is shown by the well-known gedanken example of a computerized propositional knowledge base built to 'describe' freshman calculus, but which happens to be identical to one that 'describes' Napoleon's battle at Waterloo: the semantic interpretation depends on the *user's* intent, and not on the representation at all. It is the embedding of a formal system in the world that gives meaning to the system.

It seems 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. When we want to discover whether or not we are near a door, we often postulate predicates like '**doorp**'—but we rarely worry about exactly what goes on in evaluating **doorp**, and so we don't really know what '**doorp**' means. The lack of concern with how our representations are embedded in the world appears to us in large part responsible for the problems encountered in dealing with meaning and intentional attitudes in the domains of planning, natural language, and knowledge representation.

The propositional doctrine's vision of AI as the study of logic-like propositional formalisms is doomed to failure because, by restricting itself to the study of formal languages, it cannot hope to adequately address problems of meaning. Meaning derives solely from the relationship between the world and our representational systems—the subject of scientific study.

The AI research paradigm we champion is mentioned in Nilsson's article as the "form follows function approach:" i.e., start with a functional characterization of what needs to be computed, and then figure out how to do it. More to the point, our view is that AI needs a research paradigm that addresses the problem of intelligence in terms of how meaning is obtained, extended and used—and that requires focus on making our representations capture meaning.

The Role of Logic in Artificial Intelligence Research

Although logic-like formal languages cannot be the exclusive focus of AI, they clearly will play *some* role in formulating representations and in reasoning. The 'propositional doctrine' places them centrally, indeed, it claims that such languages are the *only* reasoning tools that need be considered. There are, however, several reasons for suspecting that logic will be as inadequate for general-purpose reasoning as it is for investigating meaning.

Logic Is Not A Privileged Language

The first point to make is that logic (e.g., the predicate calculus) certainly *is* a universal language; in fact, anything that can be accurately specified can be stated in logic. We can have no objections about the formal adequacy of logic

At first glance, this seems to give the Logic Imperialists their whole case—but in fact it is a quite minor point. There are several formalisms that could be called 'universal languages,' i.e., they allow us to describe any well-formed concept and provide mechanisms for reasoning about these concepts. Obvious examples are the Turing machine and Bayesian decision theory. Each of these formalisms can describe the others, and therefore are equivalent in expressive competence. This does *not* imply, however, that they provide equally useful descriptions.

The distinction that is often lost under the statement "logic is a universal language" is the distinction between the notions of "can describe" and "is useful or appropriate to describe." We *can* describe biochemistry, cells or even human populations in terms of quantum mechanics. We don't do that because the description is too complex to be useful, and more importantly it is *inappropriate* because the description isn't in terms of the agents and objects within the domain .[2]

The point, then, is that to obtain a useful and appropriate description of a phenomenon we have to use a representation whose structure and elements are appropriate

[2]In more technical terms, the ability to derive lawful relations (those able to support counterfactuals) within a reasoning domain requires that we phrase our descriptions in terms of the natural kinds of that domain

and natural to the domain. The Logic Imperialists appear to recognize this necessity; their belief, however, is that the appropriate top-level reasoning structures will be formal, logic-like languages that use logical operations and the rules of inference to draw conclusions. The belief that logic is the natural language of reasoning, then, is the contentious point in the statement "logic is a universal language."

It is clear that logic is the natural language for explication of the proof process, or for checking the consistency of a reasoning process. This power stems from its ability to derive consistent consequences from consistent premises. When the premises are known to be inconsistent or incomplete, however, the naturalness of logic is much in doubt. Similarly, logic-like languages are naturally adapted to a discrete, state-space description of a problem, but significant contortions are required in order to deal with tasks cast most naturally in terms of continuous variables.

Deductive Versus Inductive Reasoning

It might, however, be claimed that the problems posed by biological intelligence require the discrete, consistent nature of logic-like languages. Does logic have properties that make it especially suited to the study of intelligence? We have observed that logic is well suited to proofs and deductive reasoning, both of which require consistent and sufficient initial premises. Real-world problems, however, are often characterized by both insufficient and inconsistent evidence, rendering the deductive proof process impossible.

Insufficient and inconsistent evidence is prototypical of problems in which we must "figure out what's going on," i.e., *perceptual* tasks, tasks in which we assign *meaning* to initially undifferentiated evidence. In order to deal with these problems inductive rather than deductive reasoning is required.[3] Although induction is not well understood, the one thing that *is* clear is that logical operations and the standard rules of inference are inadequate to support the inductive process. The fact that logical formalisms are competent to describe whatever inductive process we eventually discover is irrelevant *Of course* logic (and any other universal language) can model the inductive process; they are, however, *inappropriate* descriptions of the process — which is why we understand deduction but still don't understand induction.

Is Deductive Theorem-Proving Formally Adequate For Most Problems?

The deductive theorem-proving approach advocated by the propositional doctrine seems to be inadequate even for many deductive problems in which the evidence is both sufficient and consistent. Consider this example: Let us suppose that two backpackers want to divide up their load in a fair manner, by body weight. Let us further suppose that they cannot simply make a division proportional to body weight because they have very different body weights and as the division becomes more unequal the loads become proportionally more cumbersome, perhaps because of constraints on container size. The natural algebraic expression of this story problem[4] is as a sixth-order polynomial in the ratio between the carried loads.

If our backpackers proceed by putting the axioms of algebra into their theorem-proving inference engine, they will get no answer because there *is* no closed-form solution for this problem. If, on the other hand, they try to obtain an iterative approximation to the solution they will quickly obtain a useful answer—but *only* if they employ a representation that is isomorphic to the physical variables of the problem.

A representation whose structure is isomorphic to some properties of the domain being represented is able to *implicitly* represent those properties preserved by the isomorphism. We will refer to these as *isomorphic representations.* An image is a good example: distance between pixels is isomorphic to apparent distance in the world—thus *all* of the scene's metric information is implicitly represented by relations among pixels[5].

In our example the information implicitly captured by the polynomial representation is the continuous, Euclidean nature of the variables involved: weight ratios that are numerically close correspond to physically similar situations. The numerical techniques that allow solution of this problem make use of this implicit topology, and it allows them to converge on an answer.

In this example a theorem-proving approach cannot obtain a solution. We can, however, achieve success by use of a reasoning process that capitalizes on the representational isomorphism in order to construct an iterative approximation scheme. These iterative operators bear little resemblance to logical operations or the rules of inference, and yet they accomplish an important type of reasoning. It seems, then, that if we are confined to use of logical operations on our representation, as required by the propositional doctrine, then we will not be able to address this type of problem.

---

[3]Here we speak not of mathematical induction but of logical induction, i.e , the process of generating hypothesis and confirming or disconfirming them

[4]i e., find x such that $\frac{d}{dx}\left[(xw_2 - w_1)^2\left(\frac{w_1}{x-w_2}\right)^2(kx-k)^2\right] = 0$ where w_1, w_2 are the body weights of the backpackers, x is the unknown ratio of their loads, and k is the proportionality constant describing how cumbersome the load is. The solution to this equation is the simultaneously minimally-unfair and minimally-cumbersome problem solution

[5]The key idea of an isomorphic representation is that, to the extent that the constraints and relations of the problem domain are captured by the structure of the representation, these constraints and relations need not be made explicit. In some cases, such as in perception of the natural world, attempting to make such knowledge explicit would be hopeless both from a practical standpoint—because of the amount of knowledge needed—and because we have a procedural understanding of some aspects of the physical world without the corresponding propositional knowledge.

Unfortunately for the Logic Imperialists, the backpacker example is not a rare, isolated case. Problems that have trans-algebraic solutions are the rule rather than the exception, except in high-school classes and toy worlds[6].

Some will claim that we can put such isomorphism-based reasoning processes into a logic-and-theorem-proving framework by procedural attachment. But if all the actual reasoning is accomplished in procedures, does it make sense to say that we are reasoning by use of the logic? We think not. Further, consider the psychological phenomenon of mental imagery. Mental imagery is certainly iconic (visually isomorphic) in nature,[7] and it is certainly an important type of human reasoning. When so much of our interesting *cognitive* behavior is accomplished by use of our visually-isomorphic mental imagery, does it make sense to say that it should be thought of as merely procedural attachments to our logical abilities? Again, we think not.

How is it that an isomorphism-based reasoning process can succeed where the theorem-proving approach fails? The trick is that it makes use of the *semantics* of the problem domain in order to derive useful approximations. In our example the sixth-order polynomial serves as a model of the relevant aspects of the problem domain—i.e., it is a statement of the semantic considerations in this problem. The iterative solution process examines this model and constructs simple, solvable approximations to it that preserve the semantics of the original domain. That is, the approximation captures the "important aspects" of the original model and is simpler only in ways that "don't matter." The error in this approximate solution is used as an experimental observation, in order to discover how the approximation failed to capture the "important aspects" of the original model. This information is then used to construct an improved approximation.

Our point is that general-purpose deductive tools that depend only on the predicates' syntatic relationships and not on the nature of the predicates' referent (e.g., unification) are not sufficient for general-purpose reasoning. This is not to say that numerical techniques are required to solve problems such as the backpacker example: techniques other than strictly numerical ones can — and perhaps even should—be used to solve such problems. In order for a reasoning technique to reach a solution, however, it must make use of the continuous, Euclidian nature of the variables and the problem space's topology. It must, therefore, employ an explict model of the problems' semantics—one that is isomorphic to our sixth-degree polynomial—in a manner analogous to numerical methods' use of such a model.

[6] Observing the close analogy between finding algebraic (diophantine) solutions and proving logic theorems of the form $\exists x + (x)$, it seems likely that *any* purely theorem-proving approach to these problems will *usually* fail

[7] The confused and irritating debate in the psychological literature, now thankfully subsided, seems to have concluded that mental imagery is iconic, at least in the sense of being visually isomorphic.

Conclusion

Nils suggests logic as the core topic of AI, but the core topics of AI seem to be already well-defined: the study of the computational problems posed by the interrelated natural phenomena of reasoning, perception, language and learning. These phenomena may, of course, be examined by other disciplines but this does not pose a challenge to AI's viability. The vigor of our discipline will depend not on assertions about the uniqueness of our approach, but on the contributions we make to understanding intelligence and intelligent behavior.

Intelligent behavior requires rational, and not merely deductive, reasoning. There is no question that deduction and logic-like formalisms will play an important role in AI research; however, it does not seem that they are up to the Royal role that Nils suggests. This pretender King, while not naked, appears to have a limited wardrobe.

NOTICE TO ALL MEMBERS

Recently, it has come to our attention that individual(s) have been misrepresenting themselves as AAAI staff members in order to gain access to confidential information about personnel histories and salaries and corporate organizational structures. It is not the AAAI's practice to want or need such information.

We do not know who these people are or what their intentions may be. So, please be advised if such individuals contact you, please note their names, addresses, and phone numbers and confirm the intent of their call with the AAAI office (415-328-3123).

Thank you.