THE CURSE OF FREGE

(position paper for logic and rationality panel)

Ronald P. Loui
Depts. of Computer Science and Philosophy
University of Rochester
loui@cs.rochester.edu

*currently at
Dept. of Computer Science
Stanford University
loui@russell.stanford.edu

ABSTRACT

I mention two methodological problems with AI's study of inference that have to do with excessive faith in classical logic. The first is the familiar bias toward formalisms that incorporate deductive patterns of argument: a bias against inductive patterns of reasoning. The second is the pretension that inference can be studied with no regard for one's habits of representation.

By implicating Frege, I am not trying to argue an historical thesis, but I will explain why I think of Levi's phrase, "the curse of Frege," whenever I think of these problems.

POSITION

We owe so much to Frege, as logicians, that when something goes wrong with our use of logic, our identification of the problem with Frege is as much a tribute to his work as is our acknowledgement of Frege when we use logic happily. My part of this discussion on logic and rationality has to do with AI's unhappy use of logic.

Isaac Levi named a chapter of his recent epistemology tome, "The Curse of Frege," and I think of that phrase whenever I think of two methodological problems of AI that can be traced to zealous reliance on logic. I think Levi's use of the phrase can be related to my second problem, but irrespective of whether it can, I am happy to appropriate his title.

The first problem is that there is a bias of AI research toward reasoning that follows familiar deductive logical patterns, toward reasoning that mathematical and philosophical logicians would recognize. To see this, consider that any attempt to formalize non-monotonic reasoning as an addition to first order logic presupposes that rational belief formation will conform to all of first-order logic's rules for deductive inference. So if we add default rules to classical logic, and use assertion to represent rational belief in a proposition, we are already committed to consistency, and to closure, e.g., to unlimited adjunction. We cannot talk about the state of belief in which the agent believes p, believes, q, and fails to believe the conjunction of p and q: a state that some epistemologists think may be a rational state of belief (pace Kyburg).

I think of Frege when I think of this deductivist bias because of his preeminence in the formalization of this kind of reasoning.

The alternative is to dabble in the logic of rational belief, which may include principles of acceptance, confirmation, perhaps explanation and analogy, which together lead to reasoning patterns of which philosophers of science might approve. The transformations on sentences that preserve acceptability might not include all of those transformations that preserve semantic truth or satisfiability.

There is little disagreement that there has been this bias, a methodological wrong-turn through the garden, and happily as well, that this bias seems to be waning with the interest in evidential reasoning and generalization. We can think of Gil Harman's Change in View here, which takes some time to debunk the idea that rational reasoners must adopt even the most basic closure and revision principles. Of course, there have been earlier explorers of this path who have been less psychologistic. We should also think of David Israel's call to study non-monotonic reasoning without the assumption that the best system is going to be the one that most resembles first-order logic, and without pre-occupying ourselves with model-theoretic translations into set theory in the name of semantics. Finally, we can think of Peter Cheeseman's "Defense of Probability," in which he speaks of forcing the "logical straitjacket."

The second problem has to do with the interdependence of knowledge representation and inference. It should be the first page on the subject of knowledge representation, but has become important enough to mention only lately. Mostly, the emergence of this

problem is the result of alternative styles of reasoning in situations of multiple inheritance or multiple non-monotonic extensions. It is what Touretzky, Horty, and Thomason had in mind when they spoke of the "clash of intuitions." It is the main reason we cannot force the last breath out of the Yale School's proposal for chronologically ignorant non-monotonic reasoning (the postponement of abnormality until the latest possible time, e.g., Hanks and McDermott, Shoham). To appreciate the problem, the first crucial idea is that altering the meaning postulates and inference rules for our language concomitantly alters how we might write things down. The second crucial idea is that there is going to be a choice between meaning postulates and inference rules, especially for non-monotonic inference systems, where everyone seems to have a favorite way of doing things; and our evaluation of this choice must depend on how we characterize our willingness to write things down in various ways. So the problem is that until we have some way to characterize our habits of logical language use, we cannot properly criticize or vindicate our many proposals for systems of inference.

This is no more a curse of Frege than a curse of any formalist, including predecessors Boole and Leibniz, and especially implicating Carnap, a hero of the inductivists. But I think we should point to Frege again because his delivery of a compelling and powerful first-order logic resulted in the logicians' postponement of the problem of how we should match our formal language to our apprehension of the situations we want to represent. Post Frege, texts on philosophical logic had only to pay a few lines to the idea that "if . . . then . . ." was assertible in slightly different circumstances from the ordinary language conditional, and perhaps a few lines to the idea that ". . . and . . ." did not imply temporal relations of any events mentioned. In the history of logic, however, especially prior to Frege, the problem of using logical language is often as important as the design of the language.

My favorite example of this problem has to do with chaining of non-monotonic reasons.

If we write that A is reason for B, and B is reason for C, we wonder whether we can conclude C under the presupposition that A. In logics of high probability, or epsilon-high probability, such as the logic of Ernest Adams, recently reworked by Judea Pearl, we cannot guarantee that the probability of C approaches unity as the conditional probability of B upon A, the conditional of C upon B, and the probability of A, approach unity. So in these logics, C is not warranted by A. If we were in a world in which it seems that C should be warranted by A (if we were attempting to represent our apprehension of some situation in which it seemed that C is warranted if A is warranted), and if we insist on using Adams' logic, then we should write down that A is reason for C.

The alternative of course is to use an inference system in which there is chaining. Part of what it means to write down that A is reason for B is that all things being equal, A will be reason for whatever B is reason for. Part of what it means to write down that A is reason for B in such a system is that we acquiesce to chaining. When the license to chain is supposed to be withheld, it must be withheld explicitly, with a censor, such as A is reason for B, B is reason for C, and A interferes with C. Donald Nute's non-monotonic logic based on a conditional logic is a good example of a logic in which this kind of censor can be expressed.

We have a choice when we want to write down that we think we're in a situation in which A should non-monotonically yield B, but not C, yet B should non-monotonically yield C. We can use the first system with the axioms that A is reason for B, and B is reason for C. Or we can use the second system with the information that A is reason for B, B is reason for C, and A interferes with C. The choice depends on how frequently we want to censor chaining, and how frequently we want to allow chaining. It is presumptuous to allow it, but it is also extremely convenient.

The problem is that when we are moved to write down that A is reason for B, and B is reason for C, we often don't know whether we should allow C upon A. That's the whole point after all; the logic is supposed to tell us whether we want to conclude C. It is in these situations that it matters which system we choose, i.e., what language we seem to be speaking: one in which reasons chain, or one in which reasons do not. And the idea that we can figure out what language we seem to be speaking when we're moved to write down that A is reason for B, and B is reason for C, is going to occur only to those who have their hands on some kind of inductivist theory-selection mechanism.

I have an opinion on whose formal mechanism could be used here, but this is not the place to argue that opinion.

More important is the conventionalist view that emerges. There is not going to be a way of falsifying the Yale School's proposal for chronological minimization. No matter how embarrassing the exposed aberrations of chronological minimization, the entrenched believers in the principle can say that aberration is due to incorrect representation of knowledge: that the sentences used in the embarrassing examples tacitly mean something other than what we have supposed they mean, or that the principle is usually employed in different situations. It is a tautology that knowledge represented in a logical language L entails the belief or action that L says it entails. We cannot say that L is incorrect. All we can really say is that for us, the use of a particular language L is cumbersome, or that despite our best efforts, we cannot quite learn how to use it properly, or at least, that the payoff of our trying to use this language in those few glowing successes does not exceed the cost of the multitudinous errors in our unsuccessful attempts to use it. Having said this, we should perhaps decide what we would regard as progress in the design of non-monotonic (or any) inference systems.

David Poole recently made a passing distinction between communication convention and reasonable assumption. Presumably, we decide on conventions, such as whether we are talking about reasons that chain (A-reasons), or reasons that don't (B-reasons); then we attempt to use those conventions to represent empirically motivated relations, such as 'Bird x' is reason (e.g., B-reason) for 'Flies x', all things being equal. This is exactly the kind of convention at stake here. But an extreme conventionalist holds that there is no way to draw the line between convention (the analytic) and reasonable assumption (the synthetic): is it that most things we call birds fly, or that flying is part of what it means to be a bird, because if something didn't fly, we'd agree not to call it a bird? Attacks on the conventions can be protected by hedging on the defaults, and assaults on the defaults can be protected by altering the conventions.

The dilemma for designers of new logical language: writing down knowledge in new logical language can seem like radical translation, even if it's your own knowledge you're writing down. So are we inventing better systems of inference, or are we just imposing new constraints on translations?

I can't blame the conventionalist's dilemma on Frege, since conventionalism pre-dates him. But I blame the following. First, a form of logical reasoning, for which Frege is largely responsible, has stolen the show from inductivists and their comprehensive programme for rational belief. Second, the success of this form of logical reasoning seduces researchers from the truth of the conventionalist view. The consequence of a shining, rallying paradigm, such as classical first-order logic, is that it restores hope to the fascists and realists who hold that there can be a correct logic, and that its sheen of correctness can be seen from an objective view, a view from nowhere. AI, where "knowledge representation" refers to a whole subject, where choice of logic is like choice of programming language, should understand that this is not the case, that rationality cannot be guaranteed simply by demanding that some privileged language be spoken.

References

Adams, E. "Probability and the logic of conditionals," in J. Hintikka and P. Suppes, eds., Aspects of Inductive Logic, Reidel, 1966.

Cheeseman, P. "In defense of logic," IJCAI, 1985.

Hanks, S. and D. McDermott. "Default Reasoning, non-monotonic logics, and the frame problem," AAAI, 1986.

Harman, G. Change in View, MIT, 1986.

Israel, D. "What's wrong with non-monotonic logic?" AAAI, 1980.

Kneale, W., and M. Kneale. The Development of Logic, Oxford, 1962.

Kyburg, H. Probability and the Logic of Rational Belief, Wesleyan, 1961.

Levi, I. The Enterprise of Knowledge, MIT, 1980.

Nute, D. "LDR: a logic for defeasible reasoning," ACMC Research Report 01-0013, Advanced Computational Methods Center, U. Georgia, 1986.

Pearl, J. "Probabilistic semantics for inheritance hierarchies with exceptions," TR93, UCLA Cognitive Systems Lab, 1987.

Poole, D. "Defaults and conjectures: hypothetical reasoning for explanation and prediction," Report CS-87-54, U. Waterloo Logic Programming and AI Group, 1987.

Quine, W. Word and Object, Harvard, 1960.

Shoham, Y. "Chronological ignorance: time, nonmonotonicity, necessity, and causal theories," AAAI, 1986.

Touretzky, D., J. Horty, and R. Thomason. "A clash of intuitions: the current state of nonmonotonic multiple inheritance systems," IJCAI, 1987.