Sources of, and exploiting, inconsistency: preliminary report

Don Perlis
Computer Science Department
and
Institute for Advanced Computer Studies
University of Maryland
College Park MD 20742
perlis@cs.umd.edu

Abstract

Although much effort has been expended by researchers in trying to maintain a consistent belief base in formalizing commonsense reasoning, there is some evidence that the nature of commonsense reasoning itself brings inconsistencies with it. I will outline a number of sources of such inconsistencies, and discuss why they appear unavoidable. I will also suggest that, far from being a roadblock to effective commonsense, (detected) inconsistencies are often a reasoner's best guide to what to do next.

Introduction

Commonsense reasoning (CSR) is the everyday plausible reasoning we all do about (and within) the everyday world of experience. It is ordinary, usually inexpert, reasoning. It deals with natural kinds concepts, such as birds, chairs, mountains, ideas, people, cars, jobs, plans, locations, amusements. These tend not to be defined (as opposed to mathematical or legal concepts), and so easily lead to all sorts of borderline cases where confusion reigns. One might think that it is hopeless to find a consistent treatment of fallible reasoning; yet this has been one motivation for the study of nonmonotonic (NMR) formalisms. The successes of that enterprise may suggest that, despite the presence of fallibility built into the foundations of NMR, despite the inherent vagueness in many natural kinds, commonsense reasoning can (and should at its best) be consistent.

However, there are further considerations, not tied to the vagueness of natural kinds, suggesting that although the commonsense world itself (as a theory in reasoners' minds) might not be inconsistent, nevertheless the commonsense reasoning done about that world will be so, i.e., a reasoner's beliefs are almost certain to be inconsistent at times, indeed very often so. This is due to the fact that much of the data we encounter is flawed, and our reasoning must consider that. Hence we do not have axioms in the usual mathematical sense: we do not define the commonsense world we reason about: we *find* it, and we do so via ill-formed and contradictory data which we must sort through and draw plausible but temporary conclusions as we go. Even if we do get some general features right—that blocks cannot be moved unless they are 'clear'—we still cannot be sure that block B is clear; we may have mistaken block C for block B; we may have mistakenly been told block B was clear; block B may have been clear, but no longer; and so on. When we come to discover (evidence that) a belief is wrong, we are for the moment in a quandary: to believe the new evidence, or hang on to the old belief? The presence of two mutually-contradictory beliefs occasions in us the recognition of the quandary, and we treat those beliefs differently from others, until we are able to resolve the difficulty.

My examples below will not be surprising; they are everyday events. I simply hope that calling attention to them will usefully enlarge the field of concern in theoretical work in AI. Just as the observation that CSR is nonmonotonic did not scare people away from formal work but rather vastly expanded it, I believe the same can be done with respect to controlled inconsistencies.

If this is the case, if inconsistency is an essential, pervasive, and central part of commonsense reasoning, then it might be an unfortunate turn of events, for the clean sharp rays of traditional formal treatments are solidly based on consistency and model-theoretic semantics. Nevertheless, I am confident we can rise to the challenge, if need be, in constructing new and better, more robust formalisms that meet the needs of real-world agents faced with an endless barrage of conflicting information.

My own view is that inconsistencies arise in nearly all reasoning, nearly all the time, and that we are so accustomed to dealing with this in everyday life that we do not notice it, most such inconsistencies being quickly resolved. I will present a number of examples intended to illustrate this claim, and some discussion as to why these are not easily dismissed for a theory of commonsense.

I will also argue that (at least some) inconsistencies are not the great evils once thought, and on the contrary may instead partially characterize commonsense as "appropriate reasoning in the face of contradictions". As such, it lends itself to extended interactive reasoning, rather than single isolated problems where the data is fixed at the outset. Recent work (Alchourron, Gardenfors, & Makinson 1985) has bowed a bit to the idea (found in Doyle (Doyle 1979)) of studying change of belief in the face of new data, rather than inference from fixed data; but even this retains the "recency-prejudice" that newer data is better than older data—i.e., that we are given tried-and-true data to simply accept at face value, no matter how strong our evidence for earlier beliefs. Yet it is not so common—so I argue—that we are in a position to accept new data without at least a minor challenge; indeed, the challenge often helps us appreciate the new data better.

In formal dress, let RP (recency-prejudice), be the axiom

$$\alpha \in KB + \alpha$$

i.e., the arrival of new data, α, leads to an replacement of the knowledge base KB by a new version, KB + α, of which α is a member. Here it is implicit that KB, as well as its replacement, is consistent, and that fact is crucial in deciding how to shape the replacement, in which α is taken as unquestionably true.

This axiom, in one way or another, informs virtually all work in NMR, including (Alchourron, Gardenfors, & Makinson 1985; Katsuno & Mendelzon 1991). Yet it flies completely in the face of actual commonsense behavior. Often evidence for a new belief α is insufficient to warrant adopting it for long. Perhaps it is adopted temporarily while it is under consideration, or before it is realized that it conflicts with other beliefs. But to remain as a long-term belief once a conflict is noticed, α must be supplemented with evidence that it is more trustworthy than, say $\neg\alpha$. Ordinary reasoning does not come with axioms; it is a perpetually-revised network of shifting beliefs based on sifting evidence. This is not only so for humans, it must be so for intelligent robots as well, as I argue below. It is in the nature, not just of agents, but of the commonsense informational world that agents must deal with.

The idea

A new belief may contradict another already-held belief. This is not news; all of NMR is based on this idea, even though most research focuses on getting the inferences right from a fixed set of axioms, rather than on what happens when a new belief comes along.

Everyday reasoning bombards the reasoner with new beliefs, or potential beliefs, all day long. Even while passing a quiet minute resting but awake, we are often aware of time passing, of breathing, of our gaze drifting about the room, of our lazy train of thought. These awarenesses constitute beliefs about ourselves, new beliefs each moment. Thus our belief (or knowledge) base, KB, is in an almost constant state of change. A fly buzzes about me: it is here, and I believe that; then it is there and I believe that, no longer believing it is here. The new belief seems to push out the old; or does the process of seeing the fly buzz from here to there involve two parallel acts of coming to believe "there" and ceasing to believe "here"? Quite possibly in the case of some kinds of perceptual belief change, as with the buzzing fly, the new belief does not so much push the old out, as the two do a tango together of going and coming. But it is not always so simple, as examples below will show.

The new-contradicting-the-old can occur in at least two interestingly different forms: the old belief may have been in error all along but we just now discover that, or it may have been true before but now we learn the world has changed[1] In either case, the old (contradicted) belief is removed from the KB; it is removed because of the presence of the new contradictand. Thus there is a (perhaps) brief period during which both old and new contradictands are believed, and it is the recognition of their contradictory status that occasions the removal of the old. This contradiction-period may be of little importance: perhaps the new belief is kept in a potential-belief store until beliefs that contradict it are weeded out, so the KB itself might never be inconsistent. In the latter case, we might speak of an adjudication-period, during which neither the new nor the old belief is fully accepted; this period may last indefinitely, and I suggest it is the norm in CSR.

Consistency cannot in general be guaranteed; and this includes everyday reasoning, such as elementary arithmetic, and editing of text. Only when the language and inference mechanisms are severely constrained can consistency be established. Such constraints are inappropriate for CSR, since a commonsense agent has no control over what data it might receive, what new ideas it might encounter, what it might be taught, what instructions it might be given. It must rely on its own wits to decide how to treat incoming information, including, perhaps, that it does not understand the information and should for the moment distrust it.

The great preponderance of formal work in CSR assumes the belief system remains consistent; yet one of the earliest papers on the subject (Doyle 1979) outlined a theory that deliberately eschews a consistency-assumption, and indeed that specifically allows for new beliefs contradicting old beliefs. The fact that this early work was not seized upon in most later theoretical developments is, I think, an indication of the strong grip that consistency has on our theoretical imaginations. And it is a compelling vision, that there is a consistent rational belief engine with common sense.[2]

[1] Winslett (Winslett 1988) and Katsuno and Mendelzon (Katsuno & Mendelzon 1991) speak of these as requiring a revision, and an update, respectively.

[2] One many including myself have struggled to maintain (Creary 1979; Perlis 1985; 1986).

Somehow consistency carries an aura of sense, solidity, correctness, and its opposite—inconsistency—an aura of incorrectness, sloppiness, nonsense. Yet this may be an illusion. Here I hope to show that (temporary) inconsistency is not only necessary as a component of commonsense reasoning, but is even a good thing, a guide to rational decision-making.[3]

But there is another case to consider: that the new belief is wrong, and that this may even be discovered (as a plausible, nonmonotonic inference) by examining the new belief in light of many other beliefs. This case is far more common than one might think; indeed its commonplace occurrence suggests that deciding among conflicting beliefs is itself a major component of commonsense reasoning. In the next section I illustrate this thesis with a number of examples.

Most work on nonmonotonic reasoning is orthogonal to this issue, focusing mostly on getting the right (plausible) conclusions from given premises, rather than how to get from one set of premises-and-conclusions to a newer set. Yet it is this transition that reasoners are faced with, for we never have the luxury of reasoning in an unchanging world as we draw conclusions from premises. As we reason, the world changes, and our conclusions are infected with new premises (incoming beliefs) before we have finished drawing them from whatever premises we started with. To be sure, one can create artificial settings where there is no change: mathematics defines its own rules by axioms, and the name of the game is to follow those rules rather than the way of the world. Similarly, certain games have defined rules, as do blocks-world scenarios. But commonsense reasoning is reasoning in and about the real world; thus it seems incumbent on us to address the transitional aspects of reasoning. This is often called "belief change"; a longer paper will discuss approaches to belief change. Here we simply note that there seems to be almost no work on belief change that addresses the apparent need to deal with the contradiction period.

Nine examples

1. *Identification error.* A new piece of data may be plausibly reasoned to be false on the basis of old data. Consider seeing your wife drive past you in her car, and then you stare in disbelief: her car is being repaired and surely cannot be on the road. As the seen car continues along you are able to see that the license plate is not hers, and so the matter resolves in favor of your former belief: her car is indeed being repaired. In the interim, however, the KB is in a state of inconsistency, although alerted to this fact (there is a direct contradiction) and so taking care about what inferences it draws: logical consequence is *not* an adequate ground for further belief when applied to unresolved contradictands.

   Notice that in this case, we depended on another new piece of data: the license plate. However, this is not essential. Consider reading that the Hubble telescope has been repaired and is working well at last, and that it has revealed several new findings, including that the first planet, Mars, has an orbital wobble that may have led to miscalculations of its advancing perihelion important in tests of general relativity. You quickly, and without seeking further evidence, accept most of this but reject firmly the error in the name "Mars", since you are solidly convinced from many other beliefs that it is Mercury, not Mars, that the passage is about.

   This example can be altered in countless ways to suit one's taste, e.g., it can be about your local supermarket rather than telescopes and planets, but the name gotten wrong by a too-hasty reporter. You begin by believing what you are reading, since it sounds not only consistent with what you already believe, but even corroborates some of your beliefs; then you come to a statement that clashes.

   A new piece of data may eventually reasoned to be true after a period of comparison with old data; not every clash leads to rejecting the new data. Sometimes we become convinced that one or more of our old beliefs was in error; this often takes more than a single encounter, depending on how strongly we held the old beliefs in question. We may begin by fighting with new data, and in the process come to suspect some old data. In this case too, however, there is an intermediate stage of suitably reasoned inconsistency.

2. *Spatio-perceptual distortion.* We reach for our coffee cup while continuing to read. Our hand gropes and does not find the cup. We are puzzled; we find a conflict between our belief that the cup is where we are reaching, and the information from our fingers that it is not there. We experience several annoyances: at our inability to get a sip of coffee at will; at our spouse for possibly having removed the cup before we were finished with it; at ourselves for being inept at finding the cup without looking. We cannot decide whether our belief as to where we left the cup is in error, or as to where it is now, or as to where we are in fact reaching. So we look.

[3] Let me note also, however, that strict adherence to consistency also carries an aura of hidebound inflexibility, unwillingness to consider new ideas and thereby undergo (perhaps brief) periods of inconsistency. Such periods, constrained by a suitably watchful or distrustful attitude toward one's own conclusions, are familiar to everyone, and we recognize them as moments from which we often emerge with greater clarity, new insights, and re-interpretations of old beliefs. Yet we have not tried very hard to capture this "controlled inconsistency" in our theories of common sense. I believe that we will not be able to achieve programs with common sense, programs able to take advice, able to guide robots through everyday chores in everyday settings, until we discover how to let them be suitably inconsistent.

3. *Delayed discovery.* New beliefs may also be accepted outright, but only later found to be in conflict with other beliefs. Frege's published formalization of set theory was found by Russell to be inconsistent, something Frege certainly did not expect. A more mundane example is planning to take one's car in early for repairs and then to drive on to work; hopefully the truth sinks in before one is on the road.

4. *New meanings.* Familiar words often take on new meanings, as we converse with others. We find the old meaning does not fit the context, and this is often signaled to us by means of an inconsistency. You hear "he beat a tattoo on the drum" and wonder how this can be: a tattoo is a mark made under the skin. Your meaning is inconsistent with the new usage, and yet it causes you no alarm, no sense of logic coming unhinged; rather you take it right away as a new usage to learn, resolving the contradiction with great ease: this is a new usage of a word, I do not need to regard my old usage (tattoos are skin marks) as an error, except perhaps my employment of it in the present context.[4]

5. *Conflation errors.* Distinct agents may use the same word in different ways, leading to confusion. Both agents may use the name "John" but referring to two different people, and only come to suspect this after the conversation has gone on for some time. Such confusion can take the form of a contradiction: Roberta has the belief "John is tall" and hears from Stan that "John is short." She finds this to contradict her belief, and argues with Stan until they figure out that they mean two different people.[5]

   It is the contradiction between her old belief and the new data that triggers her attempt to sort things out. A robot (or person) without common sense might simply accept it all and end up believing (and doing) a lot of nonsensical things.

6. *Memory distortion.* Roberta Robot trips over a log and receives a nasty bump on her memory (KB) unit, temporarily replacing her belief B with -B. Much later she re-proves B from other beliefs, is puzzled to find the inconsistency with -B, and initiates an effort to resolve the contradiction. If her evidence for B is strong enough, she may decide -B is a mistake, and in any case she will exercise good common sense by not blandly using both B and -B for further inferencing; this requires her noticing the contradiction and taking appropriate action.

   Such behavior would be critical for an intelligent (or trustworthy) robot. One small error in memory storage could cause untold disaster if Roberta blindly accepted whatever logical consequences her beliefs lead her to. Humans constantly back off when our own reasoning tells us things we have excellent reason to distrust; and a contradiction among our inferences is a chief clue that it is time to back off and rethink.

7. *Being misled.* Roberta's problem may not be due to a fall; she may have been deliberately misled, informed that -B, and only later have time to prove B and face suitably reasoned inconsistency.

   Note that an episode of suitably reasoned inconsistency is not a state of turning off nonmonotonic reasoning; on the contrary, one needs access to one's full rational engine in order to sort out the contradiction, not to mention that yet newer data might arrive while the effort is underway. So while Roberta struggles with B and -B, C and D might arrive as new data, themselves accepted or rejected, and if accepted then perhaps used in the recovery. The case of the mistaken car above was of this sort.

8. *Outset inconsistency.* The KB may be inconsistent at the outset. This is not so outlandish. Any large-scale sophisticated KB is likely to have inconsistencies, and in any case the designers will be unable to prove its consistency. So future intelligent household robots right out of the factory crate with factory-installed start-up KB will likely have conflicting beliefs as soon as they are turned on! But these conflicts will also probably be very deeply hidden, inconsistencies but not direct contradictions that the designers would have noticed.

9. *External data distortion.* Instead of a memory problem or a lie, Roberta may have a piece of paper with the combination to a lock written on it. But it might have been recorded incorrectly, or the inkmarks might become distorted by rain. She believes it to be correct, and tries to open the lock. Upon finding that the lock does not open after several tries, she considers that the data on the paper may be wrong, which is the opposite of what she had believed. This comes about because of an inconsistency among the beliefs:

```
1. OnPaper(written-combo)
2. OpensLock(lock-combo)
3. lock-combo = written-combo
4. -OpensLock(written-combo)
```

   But it takes much more reasoning to decide that belief 3 is the culprit. For instance, Roberta may decide that that lock itself is functional, i.e., that 2 is true, but that she does not know what sequence of dialings corresponds to lock-combo, and further that it may nevertheless be similar to written-combo. She may decide to try variations on written-combo. But first she re-tries written-combo several times to be sure that it does not work; this is before she is fully sure she believes 4, although it is certainly a belief that has come to her as new data.

[4] See (Harper & Helzerman 1995) for a treatment related to this kind of phenomenon.

[5] This example is admittedly similar to identification error. See (Miller 1993; Miller & Perlis 1993; Miller 1995) for related treatments.

What kind of belief is one we are not fully sure of? It is one that we recognize as having some evidence and are taking into account, though perhaps are not ready to accept its logical consequences.

Why the problem won't go away

Why won't this issue go away? Why can't we suppose, as we suggested above for the Winslett-Katsuno-Mendelzon cases, that there is a buffer for incoming beliefs, where they are kept until the coast is clear to enter them consistently into the KB?

There are two reasons: First, we need the KB to help us decide what to believe. And the incoming belief is as much (and as little) a part of the evidence as are the other beliefs in the KB. So we must let them fight it out on equal grounds; we have no a priori grounds to favor one over the other. Indeed, if we keep the newcomer in a buffer, and as yet have not decided whether the old-timer is in error, then in effect we have put each in a buffer. Second, it is the conflict that brings about the buffer in the first place; we first need to note the contradiction before we know we need to deal with it. We cannot put all beliefs in buffers all the time just in case they might later turn out to be in conflict.

Commonsense reasoning may then, in part, be reasoning in the face of confusion due to conflicting data, reasoning when things go awry, stray from the routine, the clear, the well-behaved. I should say, *mildly* awry; for commonsense too breaks down when things get overly strained. But a mild bump to the nonmonotonic engine should not kill the engine altogether, nor rev it out of control. Flexibility in the face of (mild) straying from the middle of the road are the very stuff of commonsense; much more effort is needed in addressing this. Perhaps when we do so, we will find ourselves addressing the problem of brittleness endemic in AI systems and formalisms. One little thing out of place, and instead of reasoning "oh, this should have been such-and-such, I'll make the correction" the reasoning goes haywire, and does not even recognize that it is doing so. McCarthy's notion of elaboration tolerance (McCarthy 1993) is a step in this direction, but even that does not allow for actual false or conflicting beliefs that are detected and resolved, what might be called "perturbation tolerance". Yet the latter is the very stuff of everyday reasoning, as the above examples are intended to show.

This is not to say that our resolution of inconsistencies in our beliefs is always correct, or that we always "do the right thing" (on hindsight) with perturbations in our beliefs. In fact, even a very small change, such as believing -P rather than P, might make the world of difference in terms of whether we survive. Suppose we are told that one drives on the right in St. Thomas; this can lead to an early demise, whereas being told one does *not* drive on the right there can have a very different (better) outcome. So, the "web" of our beliefs (Quine & Ullian 1970) does not always keep us from accepting at face value a new belief, even if it is dangerous. But it does, I think, keep us from accepting it if we already have substantial evidence codified as beliefs contradicting the new one. Thus perturbation tolerance is the ability to retain a rational head in the face of modest amounts of new or altered data even if the latter contradict old data. This includes forgetting, misremembering, conflating, and other common distortions of our KB.

In the following section, I give an example of an extended episode of reasoning, in which a number of contradictions occur, and yet which a human reasoner would have no difficulty resolving, indeed perhaps would not even think that there had been anything unusual in the course of reasoning—as there would not have been: this is reasoning at its ordinary level, fraught with many small contradictions. How can a contradiction be small? It can't in ordinary logic, but there is where we need to do more theoretical work. If on December 15 I happen to come to the belief that I am 32 years old, as well as retaining the old one that I was born on January 15, 1964, then I have a minor direct contradiction in my beliefs, which I may come to discover and easily resolve.

But if my KB is full of very numerous contradictions, on all topics, direct and indirect, then it is going to be very difficult to know even how to begin to resolve this situation, for what beliefs can I use to sort out plausible from implausible? I have not suggested an algorithm for the case of minor contradiction, but I hope to have indicated why not all inconsistencies are equal; some are a reasoner's best friend, guiding us out of nonsense.

This may seem a bit at odds with our example of outset inconsistency above. The lesson may be that, if a KB has highly indirect contradictions involving many beliefs, then it is very important that that agent *not* engage in lots of deep theorem-proving, lest it become burdened with widespread contradictions it cannot hope to resolve. This may not be a problem for CSR, since it is seems to be characterized by shallow reasoning (short chains of inferences), unlike that common in mathematics.

What form a logic of perturbation tolerance might take remains to be seen. I do not think that relevance (or paraconsistent) logics are likely candidates, for they tend *not* to note and attempt to resolve contradictions; nevertheless, they do represent one possible approach to consider. Another are so-called active- or step-logics (Elgot-Drapkin & Perlis 1988). Perhaps Pollock's Oscar system (Pollock 1989) will lead to another approach as well.

Extended example

In this section I give an example of reasoning that takes place over several hours, involving many new beliefs coming in at different times, including a number that contradict existing beliefs at the time. Yet the reasoning is transparently simple, something any human could do with ease. I think that once we are able to build reasoning systems with this ability, we will be very much farther along in our quest.

While the reasoning may be far beyond what our formalisms can do today, it is important that we keep such examples before us, to help us see what directions to investigate.

Roberta and her friend Stan have the task of painting a barn today. They make the following plan: Roberta will buy the paint and brushes and Stan will buy the ladder. Stan will go to Main Street Hardware, which has a bargain on ladders, even though it is farther away than the Ellis Street paint store which also carries ladders. They expect to meet at the barn at noon and begin painting. They set off in opposite directions. Roberta heads off to Ellis Street to make her purchase. At a bridge she intends to cross she notices a sign saying the bridge is under repair and uncrossable. This forces her to back up a considerable distance and take an alternate and much less direct route. As she goes, she realizes they will need two ladders and a plank, and that she should get a second ladder and a plank at the paint store.

She then sees another sign saying that all Main Street stores are closed for the day due to a water main outage, and she assumes Stan will now try to go to the more expensive Ellis Street paint store to get a ladder. She also assumes Stan does not know about the bridge being out, since else he surely would have warned her. Roberta reasons that it will take Stan at least three extra hours to get to the bridge, notice the sign, back up to the alternate route to get to the paint store, buy the ladder, and bring it to the barn. This will make it too late to paint the barn today. She decides to purchase the paint and then wait for Stan at the barn anyway, as the only obvious place to meet up with him and replan for tomorrow; but also to forgo getting a second ladder and a plank until they can meet to work out a new plan, perhaps purchasing both ladders and plank tomorrow at Main Street Hardware.

She arrives at the paint store but the street sign reads "Ellis Avenue" rather than "Ellis Street"—she decides it must be the right place since the paint store is right there and "Avenue" and "Street" are easy to switch, and in any case she can get paint and brushes there.

She arrives back at the barn at 12:20, and finds a message from Stan, marked 11:45, saying that he has placed the ladder in the barn, and will return by 12:30 to start painting. She is puzzled by the apparent contradiction, then reasons that he must have found the Main Street hardware store open after all. Then she sees that the ladder has a tag on it reading "Main Street Hardware II, Harwood Lane location." She does not know how Stan found out about the second location of the store, but evidently he did. Since Harwood Lane is very close, she decides to make a quick trip there for a second ladder and plank. Then she reasons that she should wait instead, since it is by now 12:25 and Stan will be here at any moment.

Conclusions and future work

One's overall behavior should not become totally incoherent given a small perturbation in beliefs. This is not merely a theoretical observation; it is important for both safety and effective work toward goals, that an agent be able to behave intelligently in the face of conflicting information. It is a hallmark of common sense to be able to reason gracefully with imperfect knowledge—not only incomplete or false, but contradictory knowledge. I have illustrated this with a number of examples, and an extended episode of reasoning that any human can perform with ease, and that *not* to be able to perform would be a severe impairment of ability to get around in the everyday world. I think that until we devote serious attention to this issue, we consign out formalisms, and the robots that use them, to dull tasks not worthy of the term "commonsense".

A fuller discussion of many of these issues, and also a broader critique of where we are and where we should be, formally speaking, will be given in (Perlis 1996).

Acknowledgments

This work was supported in part by grants from the National Science Foundation and the Army Research Office.

References

Alchourron, C.; Gardenfors, P.; and Makinson, D. 1985. On the logic of theory change. *J. Symbolic Logic* 50:510–530.

Creary, L. 1979. Propositional attitudes: Fregean representation and simulative reasoning. In *Proceedings of IJCAI-79*, 176–181.

Doyle, J. 1979. A truth maintenance system. *Artificial Intelligence* 12(3):231–272.

Elgot-Drapkin, J., and Perlis, D. 1988. Reasoning situated in time I: Basic concepts. Technical Report CS-TR-2016, Department of Computer Science, University of Maryland, College Park, Maryland.

Harper, H., and Helzerman, R. 1995. Managing multiple knowledge sources in constraint-based parsing of spoken language. *Fundamenta Informaticae*. special issue on Context.

Katsuno, H., and Mendelzon, A. O. 1991. On the difference between updating a knowledge base and revising it. In *Proceedings, KR-91*, 387–394.

McCarthy, J. 1993. History of circumscription. *Artificial Intelligence* 59:23–26.

Miller, M., and Perlis, D. 1993. Presentations and this and that: logic in action. In *Proceedings of the 15th Annual Conference of the Cognitive Science Society*.

Miller, M. 1993. *A view of one's past and other aspects of reasoned change in belief.* Ph.D. Dissertation, Department of Computer Science, University of Maryland, College Park, Maryland.

Miller, M. 1995. Context shifts and clashes in dialogues: An active logic perspective. *Fundamenta Informaticae*. special issue on Context.

Perlis, D. 1985. Languages with self reference I: Foundations. *Artificial Intelligence* 25:301–322.

Perlis, D. 1986. On the consistency of commonsense reasoning. *Computational Intelligence* 2:180–190.

Perlis, D. 1996. Logic for a lifetime. in preparation.

Pollock, J. 1989. Oscar: a general theory of rationality. *Journal of Experimental and Theoretical Artificial Intelligence* 1(3):209–226.

Quine, W. V., and Ullian, J. S. 1970. *The Web of Belief*. Random House.

Winslett, M. 1988. Reasoning about action using a possible models approach. In *Proceedings, AAAI-88*, 89–93.