INTENTIONALITY AND DEFAULTS

Donald Perlis

Computer Science Department and

Institute for Advanced Computer Studies

University of Maryland

College Park, MD 20742

(301) 454-7931

perlis@cs.umd.edu

ABSTRACT

The world is too complex to always correctly model it or algorithmize responses that are always appropriate to it. For bacteria it seems not to matter, they survive in sufficient numbers without having to deal with this issue. But we are not so lucky, or rather, we are lucky that we are not so lucky, since it has forced us to evolve ways to deal with our own incorrect algorithms, namely, to postulate errors in our beliefs and, on detecting them, take corrective action. Yet to postulate error is to distinguish between error and supposed truth, and thus to recognize a sort of aboutness relation: the mistaken belief is seen as being meaningful (about something) but getting it wrong somehow. I will consider the extent to which this capacity of distinguishing between error and supposed truth may confer advantages to a commonsense reasoner, including connections to the frame problem.

I. INTRODUCTION

Mind is a device for reasoning, thinking. So, what is thought? Who needs it? Not bacteria. But more complex behavior seems to require processing information ‘about’ the world. What is ‘aboutness’ (often called ‘intentionality’) ? -- and what good is it?

The world is too complex to always correctly model it or algorithmize responses that are always appropriate to it. For bacteria it seems not to matter, they survive in sufficient numbers without having to deal with this issue. But we are not so lucky, or rather, we are lucky that we are not so lucky, since it has forced us to evolve ways to deal with our own incorrect algorithms or behaviors, namely, to postulate errors in our beliefs and, on detecting them, take corrective action. Yet to postulate error is to distinguish between error and supposed truth, and thus to recognize a sort of aboutness relation: the mistaken belief is seen as being meaningful (about something) but getting it wrong somehow.

This would seem to be essential to effective default reasoning. Thus, to hazard a very bold thought, the complexity of the world may force us to do default reasoning, which in turn may force us to have intentionality in order to sort out our mistakes. Be that as it may, I will consider the extent to which this capacity of distinguishing between error and truth may confer advantages to a commonsense reasoner.

In part II I will show how this topic bears on the frame problem. Then in III I introduce a semi-technical notion, the appearance/reality distinction (ARD), in an attempt to capture some of the above intuitions, and I provide some motivating examples from commonsense reasoning. Then in IV I quickly review the idea of intentionality as it applies to the problem of reference in the philosophy of language. This is then used in V in connection with an example of Dennett, where again the ARD appears useful. Finally, conclusions and future work are given in VI.

II. THE FRAME PROBLEM

The frame problem is I think now widely recognized as an aspect of the general problem of default reasoning. In [3], it was suggested that there are two frame problems, what we might call the 'numerical' frame problem, and the 'conceptual' frame problem. In [3] these were called the 'mathematical' and 'commonsense' frame problems, respectively. However, John McCarthy has pointed out that such terminology has the unintended connotation that the latter problem does not lend itself to mathematical formulation or solution; our intent in [3] was not this at all.

The former has to do with a precisely-defined world, in which however there is an awkwardness in expressing certain ''obvious'' facts about the inertia of most things at most times: the awkwardness is that of requiring a very great number of axioms to state this. There is nothing inherently uncertain about such a world, and default

reasoning is not needed. However, techniques from default reasoning have turned out to provide at least the beginning of a solution to the numerical frame problem. For instance, the approaches of Lifschitz [9] and Haugh [6] can be used in fully-known blocks world settings to avoid the need to axiomatize separately every contingency. Indeed, in rough form, this was Hayes' early suggestion for solving the problem [7].

Yet these approaches were born of a different intellectual concern: that of uncertainty. It was not precisely and completely characterized blocks-world settings that motivated the non-monotonic approaches of Lifschitz, Haugh, and others, but rather the more complex and uncertain world of everyday life. Of course, blocks-world scenarios were at the start intended as first steps toward the latter world. But these lacked the element of uncertainty that is now proving such a hard nut to crack. It is therefore interesting that these same approaches seem to offer hope in both settings, and (thus) for both the numerical and the conceptual frame problems.

The conceptual frame problem is as follows: The everyday world is full of uncertainties, in that it is too complex for any creature to fully axiomatize its behavior. Thus we settle for approximate rules of thumb, and (often) note them as such: birds typically fly, etc. But then to decide on the inertial tendencies of most things most of the time, we first need to have information about which things in which cases do not stay put. But there's the rub: since we don't have tried and true

axioms about this, we can only go on our rules of thumb.

Moreover, there simply are no absolutely correct rules that we might discover, because the very level of description we have chosen (birds, chairs, etc) is itself not precise. There simply is no sharp category of birds, despite zoologists' definition as feathered creatures. For a sick or very young bird may have no feathers, nor a featherless mutant that is otherwise like other birds and may even interbreed with birds. Similarly for chairs, or most concepts of everyday life. And if it is not always clear what counts as a bird, then even less can we have a clear fact of the matter as to whether they fly, or have feathers, etc. Thus the conceptual frame problem is that of accounting for stability in a world described by fundamentally vague concepts. For instance, we might conclude, as a default, that a given bird has not flown away from where we just saw it. The mentioned approaches [6,7,9] do represent gains on this, but they do not grapple with the deeper issue of inherently unclear concepts. Nor do they deal with the need to back off from such a conclusion when it is found not to correspond to facts: the bird is not found where it was. Nor again with cases in which the ill-definedness of the concept of bird confronts the reasoner directly.

In [3] we suggested that a more robust solution might involve real-time methods as in [4]. Here I want to explore a related theme, namely that backing off from a default conclusion (for instance the inertial conclusion that a bird has not flown away) as well as dealing with ill-defined concepts in general, may be

facilitated in ways that bear on intentionality to be discussed below.

To summarize the above: Most attention to default reasoning has been aimed at capturing the use of defaults, i.e., of getting the right defeasible conclusion, and not on what to do when one finds a defeater for a previous conclusion. The latter is what I have called the 'fix' problem [10]. A 'simple' solution is to throw away the defeated conclusion; but it is an unwise solution. A better one is to make defeat (and error) itself a topic about which one has extensive knowledge, for then one can explain how past errors occurred, learn to avoid similar ones in the future, and in general take account of distinctions between appearance and mention and belief on the one hand, and reality and use and truth on the other. The representation of error and the appearance/reality distinction leads directly into the above issue of aboutness.

As a point clarification, by 'appearance' I mean simply an entity recognized as essentially conceptual or mental, not necessarily viewed as corresponding to reality. Thus the reasoner may over time shift its beliefs as to what is real (true) and what is mere appearance. I am not at all suggesting that any reasoner may have the capability to make an absolute and correct judgement of what is real or true. But reasoners will, I claim, be forced by circumstances to do the best they can in this, and it is these reasoned distinctions that I am referring to.

III. THE APPEARANCE/REALITY DISTINCTION

Standard formalisms for default reasoning (including inertial reasoning as in the frame problem) do not provide conclusions about error: they either provide a default or they do not; they do not give meta-assessments of the state of one's reasoning. Here are some examples in default reasoning where explicit representation of the appearance/reality distinction (ARD) is useful in regard to error and conceptual unclarity.

A. Nixon diamond [14]: Nixon is both a Republican and a Quaker. Quakers typically are pacifists, and Republicans are not. Is Nixon then a pacifist or not? Based on the given information alone, an intuitively plausible outcome is not simply to believe nothing at all, but to realize that there are two possible appearances or conceptual outcomes (pacifist and non-pacifist) of interest and that only one of them can be true, and yet that we do not know enough to decide between them. Also we refrain from repeating the effort later -- we either leave it alone or seek more data.

B. We have a pile of seeds, which we are dropping one by one onto the ground. We release our grasp on a seed, expecting it to fall. But it sticks to our fingers. The belief that it would fall is seen to be untrue, and so we try again, this time using other means. But it may be important to remember that our first try was based on a false belief, if we now want to drop another seed. Should we proceed as we first did, or assume it too may stick? It may not matter, we can go through the whole thing again, except that this is rather slow and unintelligent behavior. It is especially vivid in the case of water on our fingers which if we notice it can lead to our

drying our fingers with a towel. But if we dry our fingers and it turns out that the pile of seeds is wet then the towel has been useless. So why did we bother drying our fingers? We cannot tell, unless we remember our thinking. And remembering one's thinking is important, as we will claim below.

C. We travel to Lower Slobbovia, and see various birds. The first one we see does not fly away when we approach quite close, contrary to our expectations. If we do not remember this, and yet continue having similar experiences, we will have no reason to alter our general expectations of birds in Lower Slobbovia. And if an unconscious weighting mechanism keeps count of defaults gone wrong (in a way that does not interact with our database and inferences) until a threshold is reached and then the default is removed altogether, then we will be in bad shape when we return home from Lower Slobbovia. That is, it is important to note that a default is not working (and even to note the circumstances). Moreover, the next time we visit Lower Slobbovia, we will want to take a special lens for close-up pictures of birds. So we need to have the revised rule explicitly represented. Finally, we end up with excellent close-ups of Slobbovian birds from all 17 of our Lower Slobbovian vacations except the first. How come? We cannot explain this except with reference to our early mistaken expectations.

This might have important repercussions. Suppose we are being questioned in court as to why we purchased a close-up lens just before our second trip to Lower Slobbovia -- it is alleged that we intended to take photos of top-secret

documents. We claim it was to photograph birds. But then why did we not buy it before our first trip? Because we learned that Lower Slobbovian birds cannot fly during our first trip, not before. Consider what the court would think if we cannot recollect this, if we actually are puzzled ourselves since we have no recollection of having had a false belief about Lower Slobbovian birds.

Another, perhaps more practical, consequence can be seen. If we want to know, for insurance purposes, what year we bought the new lens, we can figure it out by remembering that we still believed that Lower Slobbovian birds could fly until we were already on our first trip, so the lens must have been bought after that. Of course, one could simply remember perfectly when one buys things, and not need such fancy reasoning. But if we postulate a perfect memory, then surely it is odd not to allow memory of one's course of reasoning as well.

These examples illustrate that making long-term use of experiences gained in novel situations, is enhanced by having high-level access to the course of those experiences, including false starts and other errors. Still other examples can be suggested -- some of which are already in use -- such as real-time issues of taking account of where one currently is with respect to a task, which seems to hinge in part on the appearance/reality distinction (ARD) in order to separate one's goals (which are in the realm of beliefs or thoughts) from one's current state of progress (the reality).

Also, the problem-solving technique of experimenting to see what works

seems to crucially involve this same feature, especially when using a judicious mix of random trials and thought-out prospects. It makes little sense to rely purely on random trials, yet pure advance-planning is often very slow as well. A mix seems to come closer to what people do, planning a general range of likely possibilities within which to experiment, and also letting the results of the experiments re-align one's assessment of future likely possibilities. It's much like best-first search, except that it may be directly coupled to action in the environment. And then marking an experiment as evidence that something did or did not work can be very useful for future reference.

A very different realm in which ARD reasoners should excel is natural language processing. Assessing differences in usage between speakers is a canonical case of the ARD, for a word must be distinguished from its referent in order to make sense of the possibility that I am using the word in one way and you in another. That is, a word (as used by a reasoner) is an appearance, an internal or mental thing, whereas its referent, at least in many cases, is an external 'real' entity. For instance, 'John' is a word whose referent is a person, John. This observation is central to most treatments of intentionality; J. S. Mill in particular made it the focal point of his treatment of the meaning of proper names.

IV. THE WORD-WORLD CONNECTION

The word-world connection has been a traditional philosophical concern, from Aristotle to J. S. Mill, to Gottlob Frege and Bertrand Russell, to Saul Kripke and Hilary Putnam. The idea is that, somehow or other, words relate to things out there in the world, in that our beliefs, when phrased in words, lead to behavior which, if the beliefs are true, tend to get us what we want.

For a belief is useful in ordinary behavior in virtue of guiding our behavior, and to do this there must be a point at which a link occurs between verbal events and physical ones. Hence the belief 'Bush is the President' is useless unless 'Bush' is somehow tied to a person in the world. For instance, coupled with the desire to see the President and the further belief 'Bush is over there' this might lead to one's going over there. Now if 'Bush' does not have a tie to the actual Bush, then going over there and not finding Bush will not be unsatisfying, since no connection between 'Bush' and Bush is assumed. But clearly this is not how we usually construe our beliefs; we usually take them to have truth conditions, things out there in reality that make the belief either true or false (whether or not we are in a position to assess the truth of these conditions). That is, the beliefs are to us more than mere verbiage; they are 'about' the world somehow. Moreover, this explains how beliefs evolved, that is, why they are adaptive. They help us get our wants fulfilled.

Of course, many things mediate in this activity, mental events being primary among them. But we can still ask about the end connection, from word to world,

without trying to sort out all the steps along the way. The territory is complicated; for a good overview, see [2]. Here I shall attempt a very brief sketch, with apologies to philosophers of language.

Mill's direct reference theory says that, roughly, the meaning of a word is that object in the world that it stands for. By 'meaning' Mill means simply the relation between word and world, so in a sense this is a tautology. But he gets some mileage out of it anyway. He says that the truth-tie is simply that 'Bush' is the name of Bush, and thus beliefs involving the word 'Bush' are true when the corresponding property is actually held by the person Bush. This seems at once obvious and trivial, and it is. But it at least is a start toward a theory of the word-world connection.

A principal difficulty with the theory arises with multiple names for entities. Thus Bush is known to some as 'Bush', and to some others as 'George'. The sentence 'Bush is George', on Mill's account, would be trivially true since it just says that that person is himself. Yet it is clearly not trivially true; it conveys information someone might not have known. We all know a person is himself, but we do not all know that Bush is George.

Frege and Russell proposed ideas to get around this. Frege said there are two parts to meaning: the sense and the reference. Roughly speaking, the latter is the word-world tie taking a word to its correlate out there (its referent), a la Mill, and the former is the way in which the words lead us to the latter. As Russell

emphasized, they do so via a description. Thus, for a given individual, to the word ‘Bush’ might correspond the description ‘the current President of the US’ and to ‘George’ might correspond ‘the Silver Fox’s husband’. Then ‘Bush is George’ has a sense that is not trivial, for it says two very different descriptions refer to the same thing. The referential meaning of them is the same, the sense-meaning is not. One goes to find Bush by going to the White House, and to find George by going looking for the husband of the person we know as ‘the Silver Fox’. Thus one’s behavior differs in the two cases, until it is discovered that George and Bush are the same person. Behavior is based on sense-meaning.

In some ways, then this was an improvement over Mill’s view. It acknowledged that there is an important mode of presentation involved in our use of words, the sense-meaning. But it too had its problems. For instance, what description gives us the meaning of ‘Plato’? Perhaps for many people all they know is that he was the student of Socrates. But these same people might know only that Socrates was the teacher of Plato. None of this serves to make a tie between the words (‘Plato’ and ‘Socrates’) and the real people they are supposed to mean.

Enter the causal theory of Kripke and Putnam. This says that a name (‘Plato’) came into use somehow, via a particular history. Even though most of us cannot trace that history, it is nonetheless that which provides the tie between the name and the person. There really was a person named ‘Plato’ and when we say ‘Plato’ today the meaning of the word is that person named back then. (Of course

there's a lot of detail here, e.g., about borrowing the usage in good faith rather than naming your cat 'Plato' etc.) This explains how people can use a name meaningfully even if they have very little idea themselves of its world tie.

There are many subtleties I have not gone into, such as more general terms than names. General (natural kind) terms like 'water' are much harder to account for, on any of the three theories, but the causal theory plus a bit of description theory seems to do the best job.

I argue that even this will not do, however, and that we must go back and take a serious look at what goes on inside the user of the words, and not only in idiosyncratic cases where someone uses 'water' to mean 'milk' but even in conventional meanings. A bit of this idea will surface in the next section.

V. DENNETT'S 2-BITSER

I think the ARD has some signifi cance for the philosophical issues surrounding intentionality and reference. Consider Dennett's 2-bitser [1]. This is a vending machine that accepts quarters. Dennett argues -- correctly, I believe -- that the 2-bitser can be said to 'represent' or 'mean' a quarter by its internal state that results from accepting a quarter, only by virtue of an on-looker that so interprets the state. That is, the 2-bitser's intentionality is derived, not intrinsic. Dennett suggests that this is true of all intentionality, even ours. However, the 2-bitser is not an ARD device, and we are. I think that the ARD feature may lead us out of mere

derived intentionality, to intrinsic intentionality. Of course, this will in part hinge on just what we take intentionality to be. But one thing seems promising at the outset: the ARD has to it a built-in internal kind of aboutness or directedness, as Brentano originally pointed out as a requirement for intentional states. The appearance is about the reality; that is, the ARD reasoner -- by definition -- distinguishes via its representational system things that it takes to be appearances from things that it takes to be realities.

Thus the ARD requires both appearance and reality to be internally represented, for the reasoner is to explicitly reason about both, and so needs distinct tokens for the two. Both the word 'John' and the person John are to be represented, so that the reasoner can say or think that the former names the latter. (This even if John may be a figment of the reasoner's imagination, yet whom the reasoner (mistakenly) takes to exist and be named 'John'.) This already makes us different from the 2-bitser, for the vending machine has no way to regard its state as standing for anything else. That is why there is only a derived notion of representation for it: we can interpret its state as representing a quarter.

However, if the 2-bitser were equipped with a camera and suitable internal mechanisms, it could relate the state that results from accepting a supposed quarter with its visual data formed from the camera's pointing at whatever is being pushed into its slot. That is, it could treat its 'accept' state as a name for the visually parsed datum, much as we may think of the name 'John' as attached to what we see before

us (John, except of course that this is mediated by our eyes and brains, just as in the case of the 2-bitser's camera and associated mechanisms). Of course, our brains are vastly more complex than anything we currently can build into a machine, but that is another matter.

Now, one problem (of many) that surfaces here is error. Dennett points out that the 2-bitser can be fooled (with respect to its derived intentionality) by using instead of a US-quarter a quarter-balboa (identical in all signifi cant respects to a quarter, but not acceptable to the machine's owners). That is, the fooling is really with respect to people, not the 2-bitser: it is too dumb to be fooled, since it has no intentions, no interpretations of its own to be gotten around. It is people's intentions and interpretations that are fooled.

The use of a camera may offset this, by reading the inscription on the quarter-balboa and rejecting it. This might proceed by comparing the 'accept' state and the visual data and deciding that the two don't match: the former says 'US-quarter' and the latter 'quarter-balboa'. This then could cause the machine to return the quarter-balboa with a stern vocalized message to the person whoever inserted it. In colloquial terms, the 2-bitser would have caught its own mistake.

The matter will not rest there, however. For one thing, matching a canonical picture of a quarter is also not foolproof, it is no guarantee of being produced in the proper way by the U.S. Mint. Now of course, people are not very good at assessing this either. But at least we can understand the concept of being a 'real' quarter, and

recognize that this is different from our mere error-prone judgement that something is a quarter. Or so we tell ourselves. Dennett and others seem to think not, that this is an illusion about ourselves. And certainly it is difficult to say what it is that constitutes the 'real' meaning of terms, apart from our judgements.

What we would like are truth conditions for being a quarter, etc. If the conditions reside in our own judgements (verificationism) then how can we ever be wrong? And if not in our minds then where and what good do they do us? Fodor [5] and others have struggled to make good on an internal notion of error, without apparent success. What we want is to be able to be wrong and to recognize this. But to recognize it is apparently to have in mind the right answer and contrast it with the wrong. Yet if we have the right answer how do we ever come to choose the wrong one in the first place?

The causal theory of reference [8,2] tries to capture this by means of suitable generalizations based on the key causal features in the growth of a term's use, e.g., paradigmatic examples of things that came to be called quarters. This has several very difficult aspects, though it is perhaps the most robust theory around right now. One of the originators of the theory, Hilary Putnam, seems to have abandoned it in favor of the view [12] that reference is never fully and finally fixed in an external reality but rather is always relative to a language user's point of view. This would appear to be the case for our camera-equipped 2-bitser, for instance. One thing that we may have over this souped-up 2-bitser is that we can adapt our usage

as we learn more; we may start with a rather simplistic notion of quarter and then over time come to employ a far more subtle notion. For this a recursive ARD capability seems just the thing, something I call ‘reflection’ on a ‘presumed external thing’ [11]. See [13] for a related view.

Be that as it may -- and I think the ARD approach has still more to offer on this -- I think it already is apparent that there is a significant behavioral watershed when a device is able to employ the ARD. There is an internal directedness that makes symbols symbolic to the device itself, and furthermore this has behavioral adaptiveness in that a finer range of distinctions and error corrections (as in returning the quarter-balboa) becomes possible.

VI. CONCLUSIONS AND FUTURE WORK

The foregoing discussion suggests that a computational treatment of the ARD could lead to a potentially powerful tool in the AI arsenal. I have argued that it can greatly enhance the effectiveness of default reasoning, and thus aid in the conceptual frame problem, as well as other areas.

For instance, in regard to the conceptual frame problem, ill-definedness of a concept might be approached by recognizing that the concept is ill-defined, that ‘bird’, say -- although it is certainly a word (i.e., a conceptual or appearance sort of thing) -- need not correspond precisely to any reality ‘out there’ and thus it is not so surprising that it gives us occasional trouble. We then can go about trying to repair

or improve our conceptual terminology to bring it more in line with what we now take, with hindsight, to be the reality. In extreme cases we may give up our old conceptual hooks altogether (cease using the term ‘bird’) and in milder cases simply be cautious. But in any event there would seem to be the need to first realize that there is a discrepancy and that our internal model needs changing with respect to an external reality, even while both internal and external entities are actually internally represented, as we saw earlier in the example of ‘John’ and John. This is the essence of the ARD notion.

However, this has been a theoretical and philosophical discussion. A great deal of careful experimentation will be required to see whether the ARD can fulfill the promise I have suggested for it. A natural scenario for such tests would be that described in [4], and plans are underway to carry out such a study.

BIBLIOGRAPHY

(1) D. Dennett [1987] Evolution, error, and intentionality. In D. Dennett, *The Intentional Stance.* MIT Press. pp. 287-321.

(2) M. Devitt and K. Sterelny [1987] *Language and Reality.* MIT Press.

(3) J. Elgot-Drapkin, M. Miller, and D. Perlis [1987] The two frame problems. In F. Brown (Ed.) *Proceedings of the 1987 Workshop on the Frame Problem in Artificial Intelligence,* April 12-15, Lawrence, Kansas. Morgan-Kaufmann Publishers, Inc. pp. 23-28.

(4) J. Elgot-Drapkin, M. Miller, and D. Perlis [1987] Life on a desert island. In F. Brown (Ed.) *Proceedings of the 1987 Workshop on the Frame Problem in Artificial Intelligence,* April 12-15; Lawrence, Kansas. Morgan-Kaufmann Publishers, Inc. pp. 349-357.

(5) J. Fodor [1987] *Psychosemantics.* MIT Press.

(6) B. Haugh [1987] Simple causal minimizations for temporal persistence and projection. *Proceedings of the Sixth National Conference on Artificial Intelligence,* July 1987; Seattle, Washington. Morgan-Kaufmann Publishers, Inc. pp. 218-223.

(7) P. Hayes [1973] The frame problem and related problems in artifi cial intelligence. In A. Eithorn and D. Jones (Eds.) *Artificial and Human Thinking,* Jossey-Bass, Inc. pp. 45-59

(8) S. Kripke [1972] *Naming and Necessity.* Harvard University Press.

(9) V. Lifschitz [1987] Formal theories of action. In *Proceedings of the Tenth International Joint Conference on Artificial Intelligence,* August 1987; Milan, Italy. Morgan-Kaufmann Publishers, Inc. pp. 966-972.

(10) D. Perlis [1986] On the consistency of commonsense reasoning. *Computational Intelligence, 2,* pp. 180-190. Reprinted in M. Ginsberg (Ed.) *Readings in Non-Monotonic Reasoning,* 1987, Morgan-Kaufmann Publishers, Inc. pp. 56-66.

(11) D. Perlis [1987] How can a program mean? *Proceedings of the Tenth International Joint Conference on Artificial Intelligence*, August; Milan, Italy. Morgan-Kaufmann Publishers, Inc. pp. 163-166.

(12) H. Putnam [1988] *Representation and Reality.* MIT Press.

(13) W. Rapaport [1988] Syntactic semantics. In J. Fetzer (Ed.) *Aspects of Artificial Intelligence.* Kluwer. pp. 81-131.

(14) R. Reiter and G. Criscuolo [1981] On interacting defaults. *Proceedings, Seventh International Joint Conference on Artificial Intelligence, August 1981;* Vancouver, Canada. Morgan-Kaufmann, Publishers, Inc. pp. 270-276.

ACKNOWLEDGEMENTS

Supported in part by ARO grant DAAL03-88-K-0087, ONR grant N00014-82-K-0193, and NSF grant CCR-8320136. This paper was largely written while the author was on sabbatical leave at the University of Rochester. I would like to thank Elizabeth Hinkleman, Dan Dennett, Bill Gasarch, Jeff Horty, Ken Ford, and Pat Hayes for their very helpful comments.