CS-TR-3244 | March 28, 1994

Presentations and this and that: logic in action

Michael Miller and Donald Perlis*

*Institute for Advanced Computer Studies
Computer Science Department
AV Williams Bldg
University of Maryland
College Park, MD 20742

perlis@cs.umd.edu

Abstract

The tie between linguistic entities (e.g., words) and their meanings (e.g., objects in the world) is one that a reasoning agent had better know about and be able to alter when occasion demands. This has a number of important commonsense uses. The formal point, though, is that a new treatment is called for so that rational behavior via a logic can measure up to the constraint that it be able to change usage, employ new words, change meanings of old words, and so on. Here we do not offer a new logic per se; rather we borrow an existing one (step logic) and apply it to the specific issue of language change.

Note: this paper has appeared in the Proceedings of the 15th Annual Meeting of the Cognitive Science Society, held in Boulder, 1993.

Work done with partial support from NSF grant IRI-9109755.

Presentations and This and That: Logic in Action[1]

Michael Miller and Donald Perlis
Department of Computer Science
University of Maryland
College Park, MD 20742
mm@cs.umd.edu, perlis@cs.umd.edu

Abstract

The tie between linguistic entities (e.g., words) and their meanings (e.g., objects in the world) is one that a reasoning agent had better know about and be able to alter when occasion demands. This has a number of important commonsense uses. The formal point, though, is that a new treatment is called for so that rational behavior via a logic can measure up to the constraint that it be able to change usage, employ new words, change meanings of old words, and so on. Here we do not offer a new logic per se; rather we borrow an existing one (step logic) and apply it to the specific issue of language change.

Introduction

"Did you hear that John broke his leg?"
"No, really? That's a shame!"
"Yes, and his wife now has to do everything for him."
"Wife? John isn't married. Which John are you talking about?"
"I'm talking about John Jones."
"Oh, I thought you meant John Smith."

The above apparently mundane conversation hides some very tricky features facing any formal representational and inferential mechanism, whether for use in natural language processing, planning, or problem-solving. For here occurs an implicit case of language control. As it dawns on the two speakers above that they are using the name "John" differently they need to reason about usage and adopt a strategy to sort out the confusion, e.g., by using last names too.

The ability of a reasoning agent to exercise control of its own reasoning process, and in particular over its language, has been hinted at a number of times in the literature. Rieger seems to have been the first to enunciate this, in his notion of referenceability [Rieger, 1974], followed by others [Perlis, 1985], [McCarthy and Lifschitz, 1987], etc. The underlying idea, as we conceive it here, is that the tie between linguistic entities (e.g., words) and their meanings (e.g., objects in the world) is a tie that the agent had better know about and be able to alter when occasion demands. This has a number of important commonsense uses, which have been listed elsewhere [Perlis, 1991].

The formal point, though, is that a new treatment is called for so that rational behavior via a logic can measure up to the constraint that it be able to change usage, employ new words, change meanings of old words, and so on. The usual fixed language with a fixed semantics that is the stock-in-trade of AI seems inappropriate to this task.

Here we do not offer a new logic per se; rather we borrow an existing one (step-logic [Elgot-Drapkin and Perlis, 1990], [Elgot-Drapkin, 1991]) and apply it to the specific issue of language change. Referenceability, to stick with Rieger's terminology, demands that the agent – and therefore the agent's language – have expressions available to denote expressions themselves (e.g., via quotation) and also to denote the tie between an expression and what it stands for. The form that this word-object tie takes seems to vary according to context,[2] and that is what this paper will focus on, by examining several specific commonsense settings.

Traditional descriptions of nonmonotonic reasoning envision nonmonotonicity as a relationship between theories: from one theory certain theorems follow that do not follow when that theory is augmented with additional information (axioms). However, this relationship is expressed only in the meta-theory; the usual logics pay attention to behavior only within a given theory. On the other hand, "theory change" is the central feature of the *step-logic formalism*. In brief, a step-logic models belief reasoning by sanctioning inference one-step-at-a-time, where the time of reasoning is integral to the logic. Complicated reasoning made of many successive inferences in sequence take as many steps as the sequence contains. Error, change of mind, change of language, and change of language usage all are time-tempered

[1]This research was partially supported by NSF grant IRI-9109755.

[2]Recently, McCarthy and others have been investigating formal theories of context ([McCarthy, 1993], [Guha, 1991]). The implications this may have for our work are, at this point, unclear.

in that they are appropriately characterized only with regard to a historical account of beliefs, language, and its usage. The one-step-at-a-time approach offers a natural account of such histories.

A key informal idea for us will be that of a *presentation*, which means roughly a situation or context in which attention has been called to a *presumed* entity, but not necessarily an entity we have a very clear determination of at first.[3] This, we argue, is the case in virtually all situations initially, until we get our bearings. But before we actually make an identification we determine (perhaps unconsciously) that there is *something* for us to deal with. This is a small point as far as initial matters go, but becomes important if later we decide to change our usages. Some examples will help. We have devised a formalism that "solves" these example problems and have implemented our solution to some of the problems. Space allows only a brief sketch of certain underlying mechanisms.[4]

Rosalie's Car

A car flashes by us, and we quickly identify it as Rosalie's car (which for simplicity we denote rc). We may be unaware of any recognition process, thinking simply that we see rc flash by. Then we notice that the license plate on the car is not what we would expect to see on rc, and we re-assess our belief that we are seeing rc. Something, we tell ourselves, made us think *this* (the car we see driving away) is *that* (the car rc we already knew of from earlier times). Once we have produced appropriate internal tokens, we can then say that we mistook *this* for *that*. The something-or-other that brought about our mistake is what we call a presentation. It will not play a formal role for us, but simply a motivational one in leading us to our formal devices.

How can we formalize the notion of taking *this* for *that*? We begin by looking into the relationship between the two – not a physical relationship, as in features that the two cars may share (though this may ultimately have a bearing on belief revision) but rather a cognitive relationship between the entities. This relationship is suggested in the case of the mistaken car by the English statement, "I mistook *this* car to be *that* (Rosalie's)." The *this* here can be viewed as a demonstrative which (together with an appropriate demonstration) is used to pick out the mistaken car, the one which passed by. The *that* can be viewed as another demonstrative which is used to pick out rc. The statement, "I mistook this car to be Rosalie's", indicates a cognitive tie between two objects, automobiles in this case, that are in a sense linked in a (former) belief by the term rc.

Essentially what has happened is this: Initially, we are aware of an interest in one car only: Rosalie's; then later, in two: Rosalie's and the car that flashed by (i.e., the car *mistakenly identified to be* Rosalie's). In a sense, the term 'rc' in the original belief 'rc just went by' refers to both of these cars.[5] That is, we had rc in mind but connected a "mental image" of it to the wrong car, the one that flashed by. As such, beliefs about the incident reflect an unfortunate mental conflation or compression of these two cars that must be torn apart in the reasoning process.[6]

We use the 4-*ary* predicate symbol **FITB** to state that an object of perception (presented at some time or *step*) is at first identified to be some (other) object, thereby producing a (set of) belief(s), i.e., $FITB(x, y, S, i)$ says that object of perception, x, which was presented at step i, is at first identified to be y producing the beliefs in the set S. Then we use Russell's ι-operator à la Russell to pick out the *this* that was mistaken for *that*, e.g., $\iota x FITB(x, rc, \{FlashedBy(rc)\}, t)$ – "the unique object of presentation, presented at step t, which was at first identified to be rc which produced the belief $FlashedBy(rc)$." This *reality term* is used to denote what a reasoner currently takes to be some entity, possibly filling in for a previously held, but incorrect description of the same entity. (As a shorthand convention we use $tfitb(y, S, i)$, "the thing (object of presentation) which was at first identified to be ...", in place of $\iota(x)FITB(x, y, S, i)$.) By incorporating reality terms we are able to express certain errors of object misidentification reflected in one's former beliefs, for instance: $tfitb(rc, \{FlashedBy(rc)\}, t) \neq rc$ – "the unique object of presentation which was at first identified to be rc at step t, which produced the belief $FlashedBy(rc)$, is not rc. (We abbreviate assertions of the form $tfitb(t, S, i) \neq t$, (which we call *tutorials*) by $MISID(t, S, i)$.) Asserting the error sets in motion a belief revision process which is characterized, in part, by the following: The earlier belief $FlashedBy(rc)$ is disinherited,[7]

[3] The vagueness in our notion of presentation does not, at this stage, hinder our formal treatment. However, we believe it will eventually be necessary to clarify this notion. This is the focus of ongoing work. Among other things, it will involve a notion of attention, as hinted at by our informal "this" and "that" description below.

[4] See [Miller, 1993] and [Miller and Perlis, 1993] for more complete details.

[5] We assume that beliefs are symbolically represented inside the head in some mental language. [Fodor, 1979]

[6] The term *compression* is borrowed from [Maida, 1991].

[7] Disinheritance is a fundamental feature of step-logic. In particular, when two simultaneously held beliefs are in direct contradiction, neither is inherited to the next step, although either may later be re-proven by other

i.e., the step-logic ceases to have that belief, although it does retain (as a belief) the historical fact that it once had that belief, and

$$FlashedBy(tfitb(rc, \{FlashedBy(rc)\}, t))$$

is produced.

Just how does one come to suspect and detect erroneous beliefs? We have already alluded to one answer, namely that we come to suspect an error upon noting competing or incoherent beliefs. We may suspend the use of potentially problematic beliefs, perhaps speculating and hypothesizing about alternative views of the world, in an effort to hash out the difficulty. How does one decide just which alternative to have faith in? In some cases one may use a hypothesize-and-test process to ferret out the problem from the set of possible errors that might have been made. A complete principled account of how one speculates and then confirms or denies her suspicions is beyond the scope of this paper.[8] Instead, a simplifying assumption is to postulate a *tutor* or an advisor that can tell us about our errors.[9] The tutor plays the role of a friend who says, "Hey, that's not Rosalie's car". How the agent comes to represent and use the friend's advice is the issue we are addressing.

One and Two Johns

Our **One John** example is very similar to that of rc above, but will help us in moving toward the third example below. Here we imagine that we are talking to Sally about a third person, whom we initially come to identify as our friend John, merely in virtue of matching John to Sally's description of the person, or the context of the conversation, etc., but not in virtue of hearing Sally use the name "John". Later we find out it is not John, but someone else.

There is no appropriate entity before us in *perception* which has been misidentified as in the case of the mistaken car; rather it is an abstract entity, a someone-or-other, still an object of presentation, the person that Sally had in mind. There is *this* someone that has been taken to be *that*, John. Our formalism treats abstract (objects of) presentation(s) of this sort much like the case of rc.

means. Another way disinheritance allows the agent to cease believing a wff, that we introduce here, is based on a misidentification.

[8] It is likely that default reasoning is involved as is knowledge about the likelihood of errors (e.g., a car is likely to be misidentified since there are typically many similar looking cars).

[9] See [McCarthy, 1958] for a discussion about programs and advice taking.

Now let us extend this to the **Two Johns** case: We are in a situation in which we are presented with a notion of a person, whom we (come to) think is our friend John. Then we are led to believe that he has a broken leg and his wife has to do everything for him. Later we suspect that there is a confusion, that not everything we are hearing makes sense. (John, our friend, is not married.) Is Sally wrong? Or have we got the wrong person in mind? Now here is the twist: Sally starts employing the name "John" to refer to this person.[10] Perhaps she is talking about a different John. To even consider this option we need to be able to "relax" our usage so that "John" is not firmly tied to just one referent. And later when Sally says that she is talking about John Jones, not our friend, John Smith, we need a way to refer to the two entities without *using* the term $John$. We may continue to *mention* the name, but judiciously, as it is ambiguous.

We can try to employ the same formal strategy that the agent used above. Namely, we may initially come to suspect that

$$tfitb(john, BrokenLeg(john)) \neq john$$

which has the English reading: "the unique object of presentation which was at first identified to be John, producing the belief $BrokenLeg(john)$, is not John." But then once we hear Sally use the name "John" to refer to the person with the broken leg, whom we now believe is not our friend John, more must be done – the name "John" must be disambiguated.

This is where we must exhibit control over our language and language usage. First the ambiguity must be recognized. That is, we must come to see that *this* and *that* share the same name. Once that is done, new terms should be created, each to unambiguously denote one of the two Johns.

Proper naming and the use of names is made explicit with the predicate symbol **Names**. We write $Names(x, y, i)$ to state that x names object y which first came to be known (by the reasoner) at time or step i; this could be weakened to time $\leq i$, or time $\geq i$, etc., if the exact time is not known. Including the third argument is somewhat nonstandard, though not without a commonsense basis. We usually have at least a vague idea of when we come to know about someone. We can think of $Names(x, y, i)$ as collapsing $IsNamed(x, y) \wedge FirstLearnedAbout(I, y, i)$, where I is intended to be the first person pronoun.

[10] The sequence of events here is different than that reflected in the dialogue at the beginning of this abstract. Specifically, Sally uses the name "John" here only *after* we come to think that she is talking about our friend John. In the full paper we also discuss another version, in which Sally uses the name "John" at the outset.

To make ambiguity precise the binary predicate symbol **Amb** is used to state that a name does not refer uniquely beyond a certain step. Axiom **AM** expresses this:

AM : $(\forall x)(\exists yzij)\{(Names(x,y,i) \wedge Names(x,z,j) \wedge y \neq z \wedge i \leq j) \rightarrow Amb(x,j)\}$

It says that if two different objects share a name, then the name is ambiguous for the reasoner once he became aware of both objects.

Once an ambiguity arises, our reasoner will need to disambiguate any belief using the ambiguous term. We use $RTA(x,y,i)$ to state that object x is referred to as y prior to step i. In particular if $Names(x,y,j)$ then $RTA(x,y,k)$ for $k > j$, $trta(y,i)$ is used an abbreviation for:

$$\iota x RTA(x,y,i)$$

"the unique thing referred to as y prior to step i", itself a non-ambiguous reality term.

Figure 1 gives a brief sketch of the evolution of reasoning we have in mind. In the figure we use M, BL, j, and $`j$ to abbreviate *Married*, *BrokenLeg*, *john*, and *`john* respectively. Also $j1$ is used to abbreviate the expression $trta(`j,2)$, i.e.,

$$j1 = \iota x RTA(x, `j, 2)$$

namely "the unique thing referred to as 'john' prior to step 2", and $j2$ is used to abbreviate the expression $tfitb(trta(`j,2), \{M(j), B(j)\}, 2)$, i.e.,

$$j2 = \iota x FITB(x, \iota y RTA(y, `j, 2), \{M(j), B(j)\}, 2)$$

namely "the unique thing which was first identified to be the unique thing referred to as 'john' prior to step 2, which produced the beliefs $Married(john)$ and $BrokenLeg(john)$ at step 2." The predicate symbol $Contra$ indicates a contradiction between its arguments, a signal to the reasoner that something is amiss thereby initiating a belief revision process.[11]

We can view each step as a discrete moment in the reasoning process. Formulae associated with each step are intended to be (some of) those relevant to the story as time passes. At each step, underlined wffs reflect beliefs newly acquired at that step. Others, in step-logic terminology, are *inherited* from the previous step. Ellipses indicate that *all* beliefs shown in the previous step are inherited to the current step.

Beliefs at step 1 are those held before the agent's conversation with Sally and those at step 9 reflect

---

Step 1: $\underline{\neg M(j), Names(`j, j, -\infty)}, AM$

Step 2: $\ldots, \underline{BL(j), M(j)}$
(**Sally:** "...his leg is broken and his wife...")

Step 3: $AM, \underline{Names(`j, j, -\infty), Contra(\neg M(j), M(j))}$
(**Agent:** "Impossible! He isn't married.")

Step 4: $\ldots, \underline{MISID(j, \{M(j), B(j)\}, 2)}$
(**Sally:** "You misidentified who I'm talking about.")

Step 5: $AM, \underline{M(tfitb(j, \{M(j), BL(j)\}, 2)),}$
$\underline{BL(tfitb(j, \{M(j), BL(j)\}, 2))}$
(**Agent:** "So that's what's wrong.")

Step 6: $\ldots \neg M(j)$
(<Reinstate Marital Belief>)

Step 7: $\ldots, \underline{Names(`j, tfitb(j, \{M(j), BL(j)\}, 2), 2)}$
(**Sally:** "I'm talking about John.")

Step 8: $\ldots, \underline{Amb(`j, 2)}$
(**Agent:** "Oh, they have the same name!")

Step 9: $AM, \underline{\neg M(j1), M(j2), BL(j2),}$
$\underline{Names(`j, j1), Names(`j, j2),}$
$\underline{j2 \neq j1,}$
(**Agent:** "Now I've got it.")

Figure 1: Sketch of stepped-reasoning in the *Two Johns* story.

---

an unambiguous account of the two Johns, one now denoted by $j1$ and the other by $j2$, once the problem is sorted out. In between are steps whose beliefs reflect information acquired via the conversation with Sally (steps 2 and 7) and via her advice (step 4); steps whose beliefs reflect that problems have been noted (a contradiction is noted in step 3 and the ambiguity is noted in step 8); and steps reflecting disinheritance (going from step 2 to 3, and from step 5 to 6).

The indicated steps have the following intuitive gloss: (1) the agent believes that John is not married, and is named "John". Then (2) comes to believe his leg is broken and he is married. This produces a contradiction, noted in (3), so neither marital belief is retained. Advice is then taken that John has been misidentified (4) which leads to the

[11] E.g., suspending the use of potentially problematic beliefs, in particular the contradictands and their consequences. See [Elgot-Drapkin and Perlis, 1990] for details.

retraction (disinheritance) of the belief that John has a broken leg (6). The agent learns that the 'other person' is named "John" (7), notes the ambiguity (8), and takes corrective action (9) by creating and incorporating the unambiguous terms $j1$ and $j2$, one for each John.

Formal Treatment

There are several notable features of the stepped approach to reasoning illustrated in the previous section which will need to be preserved in a formal device applied to the specific issue of reasoning about former beliefs. Most conspicuous is that the reasoning be situated in a temporal context. As time progresses, a reasoner's set of currently accepted beliefs evolves. Beliefs become former beliefs by being situated in an ever changing "now", of which the reasoner is aware.

Secondly, inconsistency may arise and when it does its effect should not be disastrous; rather it should be controllable and remedial, setting in motion a fairly broad belief revision process, which includes belief retraction.

Finally, the logic itself must be specially tailored to be flexible or "active" enough to allow, even encourage, language change and usage change when necessary. As a theoretical tool the general step-logic framework developed in [Elgot-Drapkin and Perlis, 1990] and [Elgot-Drapkin, 1991] is well suited to these desiderata.

A step-logic models reasoning by describing and producing inferences (beliefs) one step at a time, where the time of reasoning is integral to the logic. Complicated reasoning made of many successive inferences in sequence take as many steps as that sequence contains. A particular step-logic is a member of a class of step-logic formalisms; each particular step-logic is characterized by its own *inference* and *observation* functions (illustrated below).

One distinguishing feature of step-logics is that only a finite number of beliefs (i.e., theorems) are held at any given discrete time, or *step*, of the reasoning process. Thus we can view each step as a discrete moment in a reasoning process.

Let α, β, and γ (with or without subscripts) be wffs of a first-order language $\mathcal{L}$ and let $i \in \mathbf{N}$. The following illustrates what a step in the modeled reasoning process of a step-logic looks like.

$$\mathbf{i}\text{: } \alpha,\ \beta,\ \gamma,\ \ldots$$

represents the belief set of the agent being modeled at step i, i.e., if it is now step (or time) i then α, β, and γ are currently believed.

A wff becomes an i-theorem (roughly, a belief a step i) in virtue of being proven (inferred) at step i. Proofs are based on a step-logic's inference function, which extends the historical sequence of beliefs one step at a time. An inference function can be viewed as a collection of inference rules which fire in parallel at each step in the reasoning process to produce the next step's theorems. For every $i \in \mathbf{N}$, the set of i-theorems are just those wffs which can be deduced from the previous step(s), each using only one application of an applicable rule of inference.

Inference rules, in their most general form, adhere to the structure suggested by rule schema **RS** below.

$$\textbf{RS:} \qquad \begin{array}{ll} \mathbf{i-j}: & \alpha_{i-j_1}, \ldots, \alpha_{i-j_m} \\ \vdots & \quad\vdots \\ \mathbf{i}: & \alpha_{i_1}, \ldots, \alpha_{i_n} \\ \hline \mathbf{i+1}: & \beta_1, \ldots, \beta_p \end{array}$$

where $i, j \in \mathbf{N}$ and $(i-j) \geq 0$. The idea behind schema **RS** is this: at any step of the reasoning process the inference of β_1 through β_p as $(i+1)$-theorems is mandated when all of α_{i-j_1} through α_{i-j_m} are $(i-j)$-theorems, and all of α_{i-j+1_1} through α_{i-j+1_r} are $(i-j+1)$-theorems, ..., and all of α_{i_1} through α_{i_n} are i-theorems.

Now we apply this to *Two Johns*. We will discuss several of the important step-logic inference rules which come into play in steps 1 through 9 of figure 1. (Others are treated fully in [Miller, 1993]).[12]

"**Observations**" can be thought of as non-logical axioms or facts which the agent acquires over time. Observations are proven in accordance with rule O:

$$\text{Rule O:} \quad \begin{array}{ll} \mathbf{i:} & \\ \hline \mathbf{i+1}: & \alpha \end{array} \qquad \text{IF } \alpha \in Obs(i+1)$$

where the function Obs is tailored to correspond to the particular problem to be solved. For *Two Johns Obs* is defined by

$$Obs(i) = \begin{cases} \neg M(j), Names(`j, j, -\infty), AM & if\ i=1 \\ M(j), B(j) & if\ i=2 \\ MISID(j, \{M(j), B(j)\}, 2) & if\ i=4 \\ Names(`j, tfitb(j, \{M(j), B(j)\}, 2), 2) & if\ i=7 \\ \emptyset & \text{otherwise} \end{cases}$$

[12] Among those not discussed here are rules for inheritance (of beliefs from one step to the next), modus ponens, contradiction handling and other belief disinheritance, and negative introspection.

which indicates beliefs which the agent held prior to "talking with Sally" (those in $Obs(1)$) and those acquired while "talking with Sally" (those in $Obs(2)$, $Obs(4)$, and $Obs(7)$). Thus the use of rule O adds new beliefs at steps 1, 2, 4 and 7 in in the solution to *Two Johns* (as depicted in figure 1).

The "**Misidentification Renaming**" rule (M) takes care of the renaming of a misidentified object in the beliefs produced by the presentation. It says this: If α, containing the term t, was produced by a presentation at step k and a misidentification of t comes to the reasoner's attention at a later step i, then at $i+1$ the reasoner will believe that α holds of the misidentified object (of presentation), i.e., $tfitb(t, S, k)$ where S is a set of wffs and $\alpha \in S$.

$$\text{Rule M:} \quad \frac{\mathbf{i}: \quad MISID(t, S, k)}{\mathbf{i+1}: \quad \alpha(t/tfitb(t, S, k))} \quad \text{WHERE } \alpha \in S$$

In figure 1 rule M applies at step 4 to produce the beliefs $M(tfitb(j, \{M(j), BL(j)\}, 2))$ and $BL(tfitb(j, \{M(j), BL(j)\}, 2))$ which appear at step 5.

The "**Ambiguity Renaming**" rule (A) disambiguates name clashes:

$$\text{Rule A:} \quad \frac{\mathbf{i}: \quad Amb(`x, k), \alpha(x)}{\mathbf{i+1}: \quad \alpha(x/trta(`x, k))}$$

This rule takes an antecedent wff $\alpha(x)$ which *uses* the ambiguous term x and eliminates the offending term replacing it with $trta(`x, k)$, which mentions but does not use $x$. In figure 1 rule A applies at step 8 to produce the beliefs $\neg M(j1)$, $M(j2)$, $BL(j2)$, $Names(`j, j1)$, $Names(`j, j2)$ and $j2 \neq j1$ which appear at step 9. (Recall that both $j1$ and $j2$ abbreviate terms which contain the subterm $trta(`j, 2)$, which is created by rule A.)

Our full system has seven additional inference rules, including a "**Name-use**" rule that can *appropriately* lead the reasoner into contradiction if names are not disambiguated.

Concluding Remarks

A formal "active" logic based on step-logic has been illustrated. This has been applied to the problem of language change, whereby an agent can alter her usage of expressions and create new expressions, as demanded by new observations.

References

[Elgot-Drapkin and Perlis, 1990] Elgot-Drapkin, J. and Perlis, D. 1990. Reasoning situated in time I: Basic concepts. *Journal of Experimental and Theoretical Artificial Intelligence*, 2(1):75–98.

[Elgot-Drapkin, 1991] Elgot-Drapkin, J. 1991. Step-logic and the three-wise-men problem. In *Proceedings of the 9th National Conference on Artificial Intelligence*, pages 412–417.

[Fodor, 1979] Fodor, J. 1979. *The Language of Thought.* Harvard University Press.

[Guha, 1991] Guha, R. V. 1991. *Contexts: A formalization and some applications.* PhD thesis, Stanford University, Palo Alto, CA.

[Maida, 1991] Maida, A. 1991. Maintaining mental models of agents who have existential misconceptions. *Artificial Intelligence*, 50(3):331–383.

[McCarthy and Lifschitz, 1987] McCarthy, J. and Lifschitz, V. 1987. Commentary on McDermott. *Computational Intelligence*, 3(3):196–197.

[McCarthy, 1958] McCarthy, J. 1958. Programs with common sense. In *Proceedings of the Symposium on the Mechanization of Thought Processes*, Teddington, England. National Physical Laboratory.

[McCarthy, 1993] McCarthy, J. 1993. Notes on formalizing context. In *Second Workshop on Formal Common Sense Reasoning.*

[Miller and Perlis, 1993] Miller, M. and Perlis, D. 1993. Language change. *In preparation.*

[Miller, 1993] Miller, M. 1993. *A view of one's past and other aspects of reasoned change in belief.* PhD thesis, Department of Computer Science, University of Maryland, College Park, Maryland.

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

[Perlis, 1991] Perlis, D. 1991. Putting one's foot in one's head – part I: Why. *Noûs*, 25:435–455. Special issue on Artificial Intelligence and Cognitive Science.

[Rieger, 1974] Rieger, C. 1974. *Conceptual Memory: A Theory and Computer Program for Processing the Meaning Content of Natural-Language Utterances.* PhD thesis, Department of Computer Science, Stanford University, Palo Alto, California.