

This preprint differs from the published version.

Do not quote or photocopy.

**CYC:  
A CASE STUDY IN  
ONTOLOGICAL  
ENGINEERING**

**ABSTRACT**

Lenat's CYC is the severest test to date of the declarative paradigm of knowledge representation and of traditional AI's 'physical symbol system hypothesis'. Lenat and Guha describe CYC as 'mankind's first foray into large-scale ontological engineering'. In designing CYC Lenat and Guha have attempted to solve traditional philosophical problems in ontology, epistemology, and logic. This paper reviews four areas that have been extensively documented by Lenat and Guha: the problem of the nature of substance; the problem of causality; the frame problem; and the problem of reasoning within an inconsistent theory or set of beliefs. The solutions to these problems that Lenat and Guha offer turn out to be disappointingly weak. Lenat's prediction that he will produce 'a system with human-level breadth and depth of knowledge' by the early years of next century betrays a failure to appreciate the sheer difficulty of the ontological, logical and epistemological problems that he has taken on.

CYC is in many ways the flagship project of mainstream AI at the present time. It represents the severest test yet of the physical symbol system hypothesis and of the declarative paradigm of knowledge representation. CYC is one of the few AI projects to take the problem of commonsense knowledge seriously. The project began at MCC in 1984 with an initial budget of US\$50 million. The goal is to build a knowledge base containing a significant percentage of the commonsense knowledge (or 'consensus reality') of a 20th century Westerner. Lenat and Guha describe CYC as 'mankind's first foray into large-scale ontological engineering' (Lenat and Guha 1990: 23). In designing CYC Lenat and Guha have attempted to solve traditional philosophical problems in ontology, epistemology, and logic.

Will the project succeed? Lenat and Guha give it nothing better than a fifty to sixty percent chance of success (Lenat and Guha 1990: 21), although given that the postulate on which the project rests - that a physically instantiated Turing machine can be organised so as to exhibit general intelligent action - is a complete unknown, there is little reason to take any figure seriously. If the project fails this may, as Guha and Lenat say, 'give us an indication about whether the symbolic paradigm is flawed and if so, how' (Guha and Lenat 1990: 57). On the other hand it may not. The project may simply collapse under the sheer difficulty of the task, leaving us little the wiser about the truth or falsity of its fundamental assumptions.

Believing as he does that at the present stage of the project the ontological, epistemological, and logical preparations are more or less complete, Lenat now places the emphasis on the mammoth task of knowledge-entry, stating that only a mere 20 percent of the project's effort need be allocated to further reflection on fundamentals (op. cit.: 192, 234). This paper reviews four areas that have been extensively documented by Lenat and Guha: the problem of the nature of substance; the problem of causality; the frame problem; and the problem of reasoning within an inconsistent theory or set of beliefs. The solutions to these problems that Lenat and Guha offer turn out to be disappointingly weak. Lenat's prediction that he will produce 'a system with human-level breadth and depth of knowledge' by the early

years of next century (op. cit.: 224) betrays a failure to appreciate the sheer difficulty of the ontological, logical and epistemological problems that he has taken on.

## 1. Ontology

Lenat records that the first five years of the project were spent primarily on the problem of devising an adequate ontology with which to underpin CYC's representational framework (Lenat and Feigenbaum 1991: 220). This not inconsiderable effort has, in my view, hardly scratched the surface. As Brian Cantwell Smith puts it, 'It's not so much that [Lenat et al.] think that ontology is already solved, as that they propose, in a relatively modest time-period, to accomplish what others spend lives on' (Smith 1991: 255). Predictably, the solutions the CYC team have adopted to ontological problems are often hasty. Here is a case study.

### *What is a Substance?*

Lenat and Guha's answer, incorporated in CYC's ontology, is that a substance is the set of all its instances. So to take a specific example, the substance gold is claimed to be the set of all the particular instances of gold that there are - particular nuggets, rings, necklaces, teeth etc. (Lenat and Guha 1990: 156ff). There are a number of difficulties with this suggestion.

(i) A person or company can own the manufacturing rights to a substance, but it is meaningless to say that one can own the manufacturing rights to a mathematical object like a set. (As Lenat and Guha somewhat inconsistently say in another context 'sets are intangible, imperceivable things; they are mathematical constructs' (Lenat and Guha 1990: 181).)

(ii) Substances can come into existence. Take the artificially created elements. Before this century there was, so far as we know, no such substance as Einsteinium. Yet sets are atemporal- they don't come into or go out of existence. (As Lenat and Guha themselves say, sets 'last forever in their Platonic universe' (Lenat and Guha 1990: 159).)

(iii) Sets can be multiplied together (we can form the so-called Cartesian product of any two sets); yet it seems meaningless to say that two substances can be multiplied together.

(iv) Some substances are more valuable than others; yet it seems meaningless to say that some sets are more valuable than others. Lenat and Guha cannot duck out of this and similar objections by insisting that 'The substance gold is valuable' really means 'Every member of the set of instances of gold is valuable'. The reason is that these two sentences have different logical implications. The sentence 'The substance gold is valuable' implies that if my buckle had been gold it would have been valuable. Since my buckle is not made of gold it is not a member of the set of instances of gold, so the sentence 'Every member of the set of instances of gold is valuable' implies nothing about my buckle. So the two sentences cannot be equivalent in meaning. I hope Lenat and Guha would shrink from the desperate measure of trying to include the "counterfactual object" my-buckle-had-it-been-gold in the set of objects that is supposed to constitute the substance gold. Otherwise one would no longer be able to affirm that every instance of gold melts at 653 °C, or has density 8.3. Why should the melting point of my-buckle-had-it-been-gold be the same as that of *actual* instances of gold? There are surely contrary-to-fact situations in which gold - that malleable, ductile, yellow metal, soluble in aqua regia and of atomic weight 179 - has a melting point of 652°C (just as there are contrary-to-fact situations in which I am an inch taller or was born with hair of a different colour).

(v) A gold nugget may contain an amount of quartz. Does this nugget count as an instance of gold? An affirmative answer leads to trouble. For by parity the nugget must also count as an instance of quartz. So the nugget is a member of both the set of all instances of gold and the set of all instances of quartz. So the substance gold and the substance quartz overlap. But this ought not to be: gold and quartz are two quite distinct substances. Indeed if it happened to be the case that *all* instances of gold contain some quartz and vice versa then the proposal would deem quartz and gold to be the *same* substance. Is the solution to say that it is not the nugget itself that counts as an instance of gold but rather the pure gold contained in the

nugget that is an instance of gold? This is to say that an instance is a distributed entity (for the gold in a lump of ore may be sparsely distributed throughout the rocky medium). How are these distributed entities to be individuated? Suppose the gold in the nugget happens to be partitioned into two regions by a thin layer of non gold bearing quartz. Do we have one instance of gold here or two? Suppose the nugget is one of several packed cheek by jowl in a seam. How many instances of gold does this family of nuggets constitute? One or many? Unless answers are given to such questions the predicate 'is an instance of gold' is not well defined, and so nor is the putative set-designator 'the set of instances of gold'. Given the ubiquity of contaminating matter, particular nuggets, rings, teeth, necklaces etc. cannot in fact be offered as paradigm examples of instances of gold by an advocate of the thesis that a substance is the set of all its instances. Yet once detached from particular rings, teeth etc. 'instance' becomes a technical term whose meaning must be explained by those wishing to use it. Until such an explanation is provided no determinate analysis of substance has been given.

(vi) The number of objects in a set - the cardinality of the set - is one of the set's *essential* properties: essential in the sense that if the number were different we would no longer have the same set but a different one. The number of instances of gold is not, however, an essential property of the substance gold. The sentence 'Gold would have been a different substance if there had been one more gold thing than there actually is' is false; yet it is true that if there were one more gold thing then the set of instances of gold would be a different set.

(vii) Sets are identical if they have the same members as one another. Suppose that (by an act of God) all the gold in the universe suddenly goes out of existence except for my gold teeth. Let us further suppose, for simplicity, that all my teeth are (pure) gold. On Lenat and Guha's proposal the substance gold = the set of instances of gold; so since it is now the case that the set of instances of gold = the set of my teeth, we get the result that the substance gold is one and the same thing as the set of my teeth. This is implausible - the more so if one reflects on the fact that in that case the sentence 'The substance gold used to have more instances than it

now does' says no more and no less than that the set of my teeth used to have more instances than it now does.

Some of these difficulties may be overcome by taking a substance to be not a set but an 'agglomeration', a single concrete yet scattered object with many parts. An agglomeration, like any other concrete object, can come into existence, be owned, and be valuable. As concrete objects agglomerations cannot be multiplied together. However, the idea that a substance is the agglomeration of its instances is in other respects no more promising than the idea that a substance is the set of its instances. The axiom of extensionality presumably still holds: two agglomerations are identical if they have the same instances. Counterfactual implications remain a problem. 'The substance gold is valuable' implies that if my buckle had been gold it would have been valuable, whereas neither 'The agglomeration of (actual) instances of gold is valuable' nor 'Each thing in the agglomeration of (actual) instances of gold is valuable' implies anything about my alloy buckle.

While the size of the KB is relatively small it is probable that CYC's shoddy ontology will not make terribly much difference to its performance. It may be predicted that as the size of the KB grows, simplifications and mistakes in CYC's ontology will disrupt the system's ability to organize and cope with the complexity of the real world.

## 2. Causality

In CYC assertions of the form 'X causes Y' are represented:

causal (X  $\rightarrow$  Y),

where ' $\rightarrow$ ' is material implication (Guha and Lenat 1990: 47; Guha 1990: 5). To take a specific example, the fact that water in the fuel causes the engine to misfire is represented like this:

causal(there is water in the fuel  $\rightarrow$  the engine misfires).

(Here I am abstracting from two small complications present in the account given by Lenat and Guha which seem merely to be matters of style, and have no bearing on the discussion that

follows. The first is that Lenat and Guha regard 'causal' as metalinguistic and prefer to record this fact explicitly in the notation by the use of quotation marks. Thus they write:  $\text{causal}('X \rightarrow Y')$ . Second, Lenat and Guha prefer to record the assertion that the embedded material implication is true by means of a separate piece of syntax. Thus they represent 'it is true that X causes Y' as:  $(X \rightarrow Y) \ \& \ \text{causal}('X \rightarrow Y')$ . (They abbreviate this whole context to '(cause X Y)' (Guha and Lenat 1990: 47).) The present treatment parallels the standard treatment of the necessity operator, L, which in terms of syntax mimics the negation operator (even though necessity is arguably a metalinguistic notion), and is such that LX entails the truth of the embedded statement X.

The symbol 'causal' is nothing more than a label or tag: the knowledge-enterers label certain implications in the KB 'causal'. CYC contains no answer at all to the problem of what it is that someone knows when they know that X causes Y. In short, CYC has no semantics for the label 'causal'.

This method of representing causality has a number of weaknesses. One of the severest of them arises from the fact that if Y is true then the implication  $X \rightarrow Y$  is true no matter whether X is true or false (and no matter whether X is even relevant to Y). (This curious property of material implication can be read off its truth table.) So both the following implications are true simply because it is true that you will, at some point, die:

You eat eggs Benedict again  $\rightarrow$  you will die.

You don't eat eggs Benedict again  $\rightarrow$  you will die.

Consider the fact that a foot descending on a french fry will cause the french fry to squash, and take some particular squashed french fry of which this causal story happens to be false: had a foot descended on it this certainly would have squashed it, but in fact what actually squashed it was Freddie's fork. Since it is true that (as Freddie's fork descends on it) the french fry squashes, the following implication is true (by the property of ' $\rightarrow$ ' just mentioned):

Dan's foot descends on the french fry  $\rightarrow$  the french fry squashes.

This is bad news for the representation

Causal(Dan's foot descends on the french fry  $\rightarrow$  the french fry squashes).

For since the descent of Dan's foot is an event of the right sort to cause the squashing of the french fry and since the implication in question happens to be true of this particular squashed french fry, someone holding a theory to the effect that there is nothing further to be taken account of in representing an episode of causation cannot resist attaching the label 'causal' to the implication. This is intolerable: if 'causal( $X \rightarrow Y$ )' is assertible when  $X$  is not the cause of  $Y$  then this formula is certainly not a satisfactory way of representing an episode of causation. The problem is particularly glaring in a case where a KB establishes 'causal( $X \rightarrow Y$ )' in this way and then reasons backwards from the truth of  $Y$  and 'causal( $X \rightarrow Y$ )' to  $X$ .

The upshot is that the labelling technique cannot be used in conjunction with material implication. Rectifying the situation is far from easy. The problem of arriving at a satisfactory formal analysis of implication that avoids the curious properties of material implication has proved to be an extraordinarily recalcitrant one.

### 3. The Temporal Projection Problem

This is a particular case of the frame problem (Hanks and McDermott 1986; McCarthy 1986). Suppose the KB contains a reasonably comprehensive collection of statements about a given situation or setting, including statements about what can cause what to happen in such a situation. We inform the KB that certain events have just occurred in the situation. To keep its description of the situation accurate the KB must be able to compute the consequences of these events. As is well known, this is far from straightforward.

Here is a toy illustration of the difficulty known as the Yale Shooting Problem (Hanks and McDermott 1986). The situation involves Freda and an assassin with a gun. The KB knows that Freda is alive when the assassin loads the gun at time  $t$ . The KB also knows that:

- (1) If a live human is shot in the head with a loaded gun they die (immediately).

The gun is not fired for two minutes. At the end of this time the assassin steps out of concealment and fires at point blank range. If told this, can the KB update its description of the situation with the assertion 'Freda is dead'?

The KB knows that:

- (2) Unless the circumstances are abnormal, if a gun is loaded at a time  $t$  then unless it is fired it will still be loaded at  $t+2$  minutes.

So if the KB assumes that the circumstances are normal (with respect to the gun's being loaded) it can infer from this 'frame axiom' that the gun is still loaded at  $t+2$  and then use (1) to conclude that Freda is dead. However, the KB also knows that:

- (3) Unless the circumstances are abnormal, a thing that is alive at  $t$  will continue to be alive at  $t+2$  minutes.

From this axiom the KB infers that Freda is alive at  $t+2$ . It can then conclude from (1) that the gun was not loaded at  $t+2$ . So the KB is in a quandary over how to update. It has arrived at two conflicting candidates: 'the gun was still loaded at  $t+2$  and Freda is dead' and 'the gun was unloaded between  $t$  and  $t+2$  and Freda is still alive'. This is the temporal projection problem: there will generally be a number of different ways of updating a description, each of which is consistent with the explicitly stated facts and each of which involves the failure of the same number of frame axioms.

Guha has proposed a solution to the temporal projection problem (Guha 1990: 7-8; Guha and Lenat 1990: 37, 46-47). The proposal makes use of CYC's representations of causality. Stripping the two candidate updates to their essentials, the first is that Freda changes state at  $t+2$  (from alive to dead) and the second is that the gun changes state between  $t$  and  $t+2$  (from loaded to unloaded). Guha points out that there is information in the story about how the first of these changes could have been caused (Freda is shot) but there is no information about how the second could have been caused. (Was it a miracle? Was the assassin torn by opposing desires?) Guha claims that CYC can readily be given the resources to exploit this asymmetry and select the first candidate. For CYC knows that

Causal(person X is shot with a loaded gun  $\rightarrow$  person X dies).

So Freda's change of state is covered by a causal statement known to CYC. CYC can check through candidate updates for changes of state that are not covered by causal statements and then select the candidate having the least number of state changes that are not covered (Guha 1990: 8).

This solution is artificial. It works only if certain causal statements are excluded from the KB. The change in the gun's state *is* covered if CYC knows that :

Causal(someone removes the bullets from a gun  $\rightarrow$  the gun is not loaded).

(It does not matter that the antecedent of this second labelled implication does not appear in the story. For nor does the antecedent of the first ('X is shot with a *loaded* gun'). The two cases are completely symmetrical.)

Withholding from CYC the information that manually removing the bullets is one way of causing a gun to be empty will certainly help CYC out of the Yale Shooting Problem, but this is hardly a recipe for a general solution to the temporal projection problem.

There is a second, deeper, difficulty with Guha's suggestion. Since CYC is a commonsense reasoner it ought to be able to conclude that the cause of the gun's emptying was not, say, activity in the LEDs in Freda's digital watch. This knowledge is represented by an assertion of the form

NOT (causal(X  $\rightarrow$  Y))

(where Y is 'the gun is not loaded' and X is a statement about the LEDs). Clearly this assertion must not count as covering the gun's change of state. Yet it does contain the label 'causal'. So when CYC checks to see what's covered and what isn't, it is not enough to look for an assertion containing the label 'causal' and an implication 'X  $\rightarrow$  Y', Y representing the event being checked. The label must be in the right position. But which position is this? Should the label perhaps be the dominant symbol? No.

(Causal(X  $\rightarrow$  Y)) AND (causal(Y  $\rightarrow$  Z))

will do to cover the occurrence of Y, but the dominant symbol is AND. Suppose we try: 'the label must either be the dominant symbol or (providing the dominant symbol isn't 'NOT') be at depth two?' No, this doesn't work either.

$$\text{NOT} ((\text{causal}(X \rightarrow Y)) \text{ OR } (\text{NOT} (\text{causal}(X \rightarrow Z))))$$

would serve to cover Z (since it entails 'causal(X → Z)'). And so it goes. There is in fact no 'right position' for CYC to look for. The method of labelling affords no easy way of establishing whether or not the KB knows a possible causal explanation for a given event.

If labels must be used then the idea of a positionality criterion should be abandoned. The obvious thing to try instead is an entailment criterion. That is, the KB should check whether what it knows entails a statement of a simple form that is certified to cover the event in question. For example, the KB should attempt to prove a theorem of the form '(EP)(causal(P → Y))' ('E' is the existential quantifier); if a proof is forthcoming, Y is covered. Unfortunately CYC will have trouble applying such a criterion. CYC may know that

$$(\text{causal}(X \rightarrow Y)) \text{ OR } (\text{causal}(Z \rightarrow Y))$$

and yet be unable to prove

$$(\text{EP})(\text{causal}(P \rightarrow Y)),$$

for CYC's general inference mechanism is only Horn-complete (Guha and Lenat 1990: 42). This means that CYC will often be unable to draw inferences that are, from our point of view, obviously correct. It is only when the premisses and conclusion of an inference are so-called Horn clauses that CYC is guaranteed to be able to make the inference. (A Horn clause is a statement whose disjunctive normal form has no more than one immediate subformula that lacks the prefix NOT.)

In summary, what the KB needs is not a bunch of ad hoc criteria based on the syntactic technique of labelling but a genuine *semantics* of causal discourse. None is in the offing.

#### 4. Reasoning the Presence of Inconsistency

In classical first-order predicate calculus any conclusion whatsoever follows from an inconsistent set of premisses. The semantical proof of this principle is straightforward. A sentence  $X$  follows from a set of sentences  $S$  if and only if there is no possible world in which all the members of  $S$  are true and  $X$  is false (this is the fundamental definition of entailment). Suppose  $S$  is inconsistent. Then there is no possible world in which all the members of  $S$  are true. *A fortiori* there is no possible world in which all the members of  $S$  are true and  $X$  is false, where  $X$  is any sentence whatsoever. Therefore,  $X$  follows from  $S$ .

Certain nonclassical logics lack this principle. The problem, though, with rejecting the principle (or any other classical principle) is that theorems cannot be excised singly from the logic. They come away in clusters and the result may be an inferential apparatus that is intolerably weak. Logics lacking the principle in question typically also lack either the disjunctive syllogism ( $\{(X \text{ OR } Y), \text{ NOT } X\}$  entails  $Y$ ) or the transitivity of implication.

Suppose CYC's 'Contradiction Detection and Resolution Module' detects an inconsistency - what can CYC do about it? CYC's first line of attack is its truth maintenance system (Lenat and Guha 1990: 50-52; Guha 1990). CYC backs up to the premisses used to derive the two conflicting assertions and, if it can, rejects one of them. (In doing this CYC makes use of information supplied by the knowledge-enterers concerning such matters as which of the premisses are least likely to turn out false, which have known exceptions, and so forth (Lenat and Guha 1990: 307).) This will not always work: CYC will often be unable to tell which premiss is to be rejected. The second line of attack is to rank the two inferences according to a number of 'preference criteria' and reject the one that scores less points. This approach is only as good as the criteria themselves and unfortunately the preference criteria that CYC contains are all rather weak. Some examples are:

(i) Prefer the inference with the stronger 'causal flavour' (Guha and Lenat 1990: 37; Guha 1990: 5).

This criterion is supposed to discriminate between inferences such as the ones involved in the Yale Shooting Problem. As we have seen, it does not work.

(ii) Prefer constructive arguments to nonconstructive ones (Guha and Lenat 1990: 37).

It is hard to see why Lenat and Guha think this criterion will be of any help. Constructive arguments are the subject matter of intuitionistic logic, and a retreat to this particular weaker-than-classical logic is of no help in dealing with the present problem, since the principle under discussion is derivable in intuitionistic logic (Gentzen 1934).

(iii) Prefer the inference with the 'shorter inferential distance' (Guha and Lenat 1990: 37; Guha 1990: 5).

This criterion will succeed in only a limited range of cases. In general there is no reason why a 'shorter' inference is more likely to have a true conclusion than a 'longer' one. The most general way of defining 'inferential distance' is in terms of the least number of steps the inference involves (for example, given modus ponens it requires two steps to infer C from the three premisses  $A \rightarrow B$ ,  $B \rightarrow C$  and A). It is clearly false that the fewer steps an inference contains, the more likely it is to lead to a true conclusion. For example, there is an inference only two steps long from the two premisses A and NOT A to the conclusion that the world ended ten minutes ago. The argument uses two inference rules.

R1: From X infer NOT ((NOT X) AND (NOT Y)).

R2: From NOT X and NOT ((NOT X) AND (NOT Y)) infer Y.

Let 'W' symbolise 'The world ended 10 minutes ago'. R1 yields

NOT ((NOT A) AND (NOT W))

from A; R2 yields W from this intermediate conclusion and NOT A.

Not every formula in the forgoing derivation is a Horn clause. However, a restriction to Horn clauses provides no help with the problems engendered by quodlibet. In particular

NOT ((NOT A) AND (NOT W))

is not a Horn clause, since its disjunctive normal form is A OR W, but this can be remedied by applying R1 to the premiss NOT A to obtain

NOT ((NOT NOT A) AND (NOT W))

and then applying R2 to this formula and NOT NOT A to obtain W. (The derivation is now three steps long, of course, since an extra move is required to derive NOT NOT A from the premiss A.) Every formula in the modified derivation is a Horn clause. In other words an inconsistent KB will be able to construct Horn clause derivations of every statement that it has the resources to express (simply substitute for W). Moreover each of these derivations will be very short - at most three steps long.

CYC's truth maintenance system and preference criteria will from time to time leave CYC unable to tell which member of a pair of conflicting assertions it should reject. At this point CYC can do one of three things.

- (1) Shut down and await human intervention.
- (2) Quarantine all the assertions implicated in the inconsistency and try to get by without them. One way of doing this is to mark each implicated assertion as being of unknown truth value (Lenat and Guha 1990: 304-305). If the number of assertions implicated is large - as it will be if one (or both) of the conflicting assertions has figured as a premiss in a large number of inferences, or if a large number of premisses were used to derive the assertions that conflict - then this may produce a considerable impairment in performance.
- (3) Brazen it out. Carry on and hope the malign effects of the inconsistency will not spread too far. This seems to be what Lenat and Feigenbaum recommend:

There is no need - and probably not even any possibility - of achieving a global consistent unification of . . . [a] very large KB. . . We expect . . . that . . . [the view that] inconsistencies may exist for a short period but. . . are errors and must be tracked down and corrected . . . is just an idealized, simplified view of what will be required for intelligent systems . . . How should the system cope with inconsistency? View the knowledge space, and hence the KB, not as one rigid body, but as rather as a set of independently supported buttes. . . [A]s with large office buildings, independent supports should make it easier for the

whole structure to weather tremors such as local anomalies . . .  
 [I]nferring an inconsistency is only slightly more serious than the usual  
 sort of 'dead end' a searcher runs into . . . (Lenat and Feigenbaum  
 1991: 217, 222)

This looks like quicksand to me. CYC is not a collection of expert-system KBs working independently of one another but a single, integrated whole. The results of inferences drawn in one butte will be available to all other buttes - and something that is available may be made use of. (As Lenat and Feigenbaum remark in a different context 'Far-flung knowledge . . . *can* be useful' (op. cit.: 199; their italics).) There is no a priori reason to think that the effects of inconsistency will not spread from butte to butte, poisoning the entire KB. In my view it would be downright dangerous to allow a commissioned KB with logic and control systems as crude as CYC's to continue to run once an irremediable inconsistency develops. For within the poisoned area, the KB will answer Yes to anything. Much damage may be done before the inconsistency is discovered (think of missile systems).

## REFERENCES

- Blair, P., Guha, R.V., Pratt, W. 1992. 'Microtheories: An Ontological Engineer's Guide'. MCC Technical Report Number CYC-050-92. Austin, Texas: Microelectronics and Computer Technology Corporation.
- Gentzen, G. 1934. 'Investigations Into Logical Deduction'. In Szabo, M.E. (ed.) 1969. *The Collected Papers of Gerhard Gentzen*. Amsterdam: North-Holland, pp.68-131.
- Guha, R.V. 1990. 'The Representation of Defaults in CYC'. MCC Technical Report Number ACT-CYC-083-90. Austin, Texas: Microelectronics and Computer Technology Corporation.
- Guha, R.V., Lenat, D.B. 1990. 'CYC: A Mid-Term Report'. *AI Magazine*, Fall, pp.32-59.
- Hanks, S., McDermott, D. 1986. 'Default Reasoning, Nonmonotonic Logics, and the Frame Problem'. *Proceedings of the Fifth National Conference on Artificial Intelligence*, pp.328-333.
- Lenat, D.B., Feigenbaum, E.A. 1991a. 'On the Thresholds of Knowledge'. *Artificial Intelligence*, 47, pp.185-250.
- Lenat, D.B., Guha, R.V. 1990. *Building Large Knowledge-Based Systems: Representation and Inference in the Cyc Project*. Reading, Mass.: Addison-Wesley.
- McCarthy, J. 1986. 'Applications of Circumscription to Formalising Common-Sense Knowledge'. *Artificial Intelligence*, 28, pp.89-116.
- Smith, B.C. 1991. 'The Owl and the Electric Encyclopaedia'. *Artificial Intelligence*, 47, pp.251-288.

