86-34

Natural Language Understanding by Computer

The Next Step

Mark S. Goldfain

The University of Illinois at Urbana-Champaign

November 1986

Abstract

A critical review of selected current research work in natural language understanding by computer is given, directing attention to some of the limitations of these working systems. Then some suggestions for possible "next" research goals are presented, followed by a discussion of two suggestions to facilitate meeting these goals.

Keywords

Natural Language Processing, Knowledge Representation, Cognitive Science, Conceptual Modeling, Language Acquisition

Natural Language Understanding by Computer ... The Next Step

§ 1. INTRODUCTION

A major thrust of current research in artificial intelligence (AI) attempts to model human skills in natural language understanding (NLU). Of course it is not entirely correct to simply label this work a subtopic of AI. In fact, it draws input from at least the five fields of philosophy, psychology, linguistics, computer science, and electrical engineering. We will briefly survey some of the established research in computerized NLU in order to highlight some of the limitations of the present state of the art. This criticism will be a bit unfair in the sense that most of this work was never intended to address issues from all of these fields at once. However, the author's opinion is that it will only be through the synergy of the newly developing interdisciplinary collaborations in these fields that we may arrive at a satisfying basic model of human language. Insofar as language is a window of the mind, our progress in understanding human language skills will go hand in hand with our progress toward understanding our thinking processes in general.

The study of language goes back at least as far as the Greek philosophers. It is actually astonishing how many of Aristotle's observations about thought and language seem to hit the mark precisely today, after more than two millenia. (See [Aristotle].) And the study of language, which is inextricably intertwined with the study of thought itself, has been mostly the realm of philosophers from that time until recently. Thus, throughout those millenia, one finds observations in the work of Hobbes, DesCartes, Locke, Frege, Kant, Hume, or Leibniz, which capture important aspects of human language.

All of this work is in some sense prescientific, however. That is, no matter how correct or important were these observations, they will have to be re-discovered in our new setting. It is one achievement (and a great one) for Aristotle to have defined the concept of reasoning by analogy, but it is yet another great accomplishment to define it with such precision that one has a computer program which is adept at performing analogies with abilities approaching those of humans ([Gentner 1983]). So we can learn much from these philosophers, but we will need to do better, as well.

In this century, this "more detailed" analysis has begun. The most notable perhaps up until 1950 would be the work of Ludwig Wittgenstein (see, for example [Wittgenstein1953] and [Wittgenstein1974], both published posthumously). These "notebook style" works do not present any cohesive theory of language, but they contain many important observations about thought, language, and the interaction between the two. A number of recent AI projects have implemented ideas that Wittgenstein put forth. (Stephen Hose Hanson, at BellCore - project "WITT". Papers forthcoming.)

Noam Chomsky has exerted tremendous influence on the field since the 1950's due to his revolutionary and pioneering contributions. (See [Chomsky1957], [Chomsky1959], and [Chomsky1965], for starters.) These give the first successful coherent attempts to formalize syntactic theory in such a way that computer implementation was possible. Chomsky argued for the *autonomy of syntax*, the belief that language understanding can be divided into two largely distinct components: a first stage, which involves *formal syntax*, and a second stage, which derives the *semantics*, or meaning, using the output of the syntactic first stage.

§ 2. RECENT WORK - IMPLEMENTATIONS

The natural language computing work of the 1960's is most noted for its mistakes and the realizations and understanding that arose from them. We will skip that period and jump into the work of the 1970's which has continued through the 1980's. Implemented systems for natural language processing can be classified in terms of the two aspects which Chomsky separated. First, the basic *parsing method* used (this being the syntactic component) and second, the underlying *representational scheme*. We take a hard look at some implemented work deriving from two major schools of thought. We will consider the work in *computational linguisitcs*, using SHRDLU as a representative example, and then we will look at *conceptual dependency* work, focusing on the work at Yale headed by Roger Schank. In both cases, our goal is to point out limitations.

§ 2.1 The Computational Linguistics School

Terry Winograd, whose very successful projects have advanced the state of the art a great deal, has largely embraced Chomsky's separation in his work to date. His SHRDLU system reported in [Winograd1971] created quite a stir in the AI community and beyond. This system involved a program which worked with a simulated table with several blocks and other similar objects. A human user would type English input and the program would respond with answers, questions, or even actions in the simulated environment. By limiting the discourse to this *blocks world*, Winograd was able to keep his program's dictionary to a manageable size and his program was able to function with a simple *procedural* representation of its understanding. Doing so freed him to handle an impressively large subset of English grammar. The major focus was on syntactic parsing of the input into constituent structures similar to those that Chomsky and others had described. (Another notable influence for Winograd was M.A.K. Halliday, see [Halliday1961].)

SHRDLU's major limitations are perhaps best pointed out by Winograd himself in [Schank1973] pp. 183-186. An excerpt follows :

> *Looking into the specific capabilities of the system, we can find many places where the details seem inadequate, or whole areas are missing. The program does not attempt to handle hypothetical or counterfactual statements; it only accepts a limited range of declarative information, it cannot talk about verbal acts, and the treatment of "the" is not as general as in the description above, and so on. These deficiencies, however, seem to be more a matter of what has been tackled so far, rather than calling into question the underlying model.*

As he continues, he makes two significant observations :

> *Looking deeper, we can find two basic ways in which it seems an inadequate model of human language use. The first is the way in which the process is directed, and the second is concerned with the interaction of the context of the conversation and the understanding of its content.*

He goes on to expand these two points. The full quotation is rather long, but to summarize, he describes the first of his two "basic" deficiencies by noting the controversy over whether a syntactic or semantic portion should control the understanding process. He suggests that better methods of shared control may be the best answer. (We will return to this issue in section 2.3) The second basic deficiency is also elaborated and here he clearly recognizes the weakness of his system in performing deductions about the usual nature of human conversation. For example, there is no attempt to model the intelligence of the partner in the dialogue, so it is impossible to make certain deductions which people commonly make and use concerning dialogue. (See [Grice 1975], [Sperber and Wilson 1986].)

His remarks are perceptive, showing that he understood quite well where his work stood and where it was headed. (Winograd's work since has continued to be very much state-of-the-art.) There are other important limitations which should also be mentioned. The system made no allowance for certain types of learning. It had limited capabilities for learning new words which were properly defined, but it could not learn anything from context, or learn new concepts, nor could it add to its understanding of grammar. This remains true of all subsequent work along this line, that the language understanding is entered by programming and is more or less "frozen". Further limitations appear below, when we deal with the state of the art in general.

A summary of further research in this tradition is now partially available in a text, see [Winograd1983]. He has continued the separation of syntax from semantics in this work to date, indeed he chose to divide his summary of the field into two books. The work *Language as a Cognitive Process* consists of *Volume One: Syntax* which is already in print, while *Volume Two: Semantics* is still in preparation.

To give a few names associated with this work, it derives from the Systemic Grammars, Generative Grammars, and the Transformational Grammars of the Chomsky school. Implementations have concentrated on syntactic parsers, devised using Transition Networks (TNs), Recursive Transition Networks (RTNs), and Augmented Transition Networks (ATNs). In terms of other details of implementation, some have developed chart parsers and others backtracking parsers.

§ 2.2 The Conceptual Dependency School

One objection to the *syntax-is-primary* approach, as Roger Schank put it, is that "the inference problem is harder than the parsing problem." The claim here is that what we do with the information in a discourse is more significant a study than how we extract it. In addition, it is known that syntax cannot always be determined without some semantic knowledge, as is evidenced by selectional restrictions, problems of finding pronomial and anaphoric referents, and so on. Schank's group has chosen to focus their research in a way which makes semantics primary.

Schank also correctly noted several psychological observations which showed human natural language use to be quite different than the results thus far achieved by programs which allow syntax to be primary or autonomous. For example, there are the failures of the generative grammar theories in predicting sentence understanding times by people. More interestingly, there is the fact that people obviously interpret sentences *as they scan them*, such that one usually has expectations that develop as one proceeds, before one has finished reading a sentence. (See pp. 12-13 [Dyer 1983].) This also occurs as we work through longer discourses, so that if we have as much as a paragraph of previous context, then we carry a great many expectations for the next sentence, which assist us in our attempt to decipher it.

A number of investigations have been carried on by Schank and his students at Yale over the past two decades. Most notably, they developed an underlying representation system for their work, called *conceptual dependency representation* (CD). This general idea has proved to be an important realization. These CD forms began as an *ad hoc* set of primitives to handle a certain domain of discourse, but it appears that the Yale group eventually realized that these were considerably more applicable than expected. In addition to building and working with CD structures, this group brought forth the ideas of *scripts* and *schemas*, and recently *MOPS*, (See [Schank 1982] and [Dyer 1983]) which can account for certain observable facts: both our building of expectations and our abilities to understand stories given very sparse information in familiar situations. The Yale Group has also begun to work on questions of learning, primarily of two types: learning the information that is understood from stories (and making further inferences and generalizations from it) and learning new scripts for story understanding (or at least revising old scripts).

A guided tour of the succession of programs they have developed over the years is given by Schank in *The Cognitive Computer* (see chapter 7, in [Schank1984]). Although these programs have made significant contributions in their intended ways, they have all had a common characteristic of being extremely limited in their vocabularies *and* their domains of discourse. These programs exhibit an "understanding" that is more psychologically interesting than previous systems. While much more can be done with their CD representations than by those systems which lacked it, still I have never yet seen an understanding system with the breadth that is found in the ATN systems of the computational linguists. It seems that current parsing technology is ahead of current representation abilities as represented by the programs we have seen thus far from the Yale school. [Note: the Yale researchers do not claim that their programs are psychologically valid, except in a very shallow way.]

§ 2.3 Uniform Limitations of the State-of-the-Art

It is possible to combine our two major approaches, merging a system which deals with CD representations of meaning on a "deep level" with one which produces the CD by means of a powerful ATN chart parser. The best merging in fact seems to be via a loose program control scheme that has been called a *blackboard*. The general idea behind a blackboard is that various independent program parts can post information as they learn it to a central data area - the blackboard. (Here the independent parts would be the ATN parser, the CD structure "selection and filling" routines, plus script knowledge that is invoked from items in the discourse.) Then each part can try to use that information, draw further conclusions, and post them, until hopefully a final answer emerges (in this case, the "most likely" parse and meaning representation for the sentence).

Such a merger has been suggested (by Gerry DeJong, for example) as a natural step. It is likely that such systems are in production at this very moment. Certainly such a system will be at least as capable as the sum of its parts. There is every reason to believe that quite good translation systems could be built in this way, for restricted domains. Such a system requires a parser for the input language which builds a CD structure, then a generator from CD to the target language (which tends to be much easier than parsing, at least to get the meaning across, if not to say it in the best way possible). To abet such efforts, there has been work on developing a transformational grammar in which all of the transformations are commutative. See Katz and Winston's paper, [Katz1982]. This program works on a large grammar and produces semantic networks as its underlying representation. Since these are formally much like CD in expressiveness, we are well on our way to this goal. At least we could shortly see good translation systems for certain restricted and well-researched domains of discourse.

All of this begs some important questions however, which have been under debate since CD, frames, and semantic networks were first introduced: nobody has successfully used these to model a significantly complex discourse using a large vocabulary. Will the methods handle it, or will they collapse either from computational complexity or due to inability to capture a wide enough variety of meanings? If they *can* handle large vocabularies in principle, will it nevertheless prove prohibitively difficult to *teach* them the large vocabulary? Will we have to construct CD or semantic nets for thousands of words by hand, with often many senses for each verb, plus procedures for them to disambiguate to the correct one? Will this involve too many man-years to complete, becoming more difficult as it grows in size?

These are some of the unanswered questions regarding the underlying representation methods in use today. Charniak and McDermott in their text [Charniak1984] (see page 8) stress the importance of representation, pointing out that an AI program in general is only as good as its data structures will allow. In the same text (see section 6.2, pages 325-333), a somewhat systematic analysis of CD representation is given. Certainly the basic CD is not general enough to handle the bulk of English without a major overhaul and extension. The particular primitives chosen do not satisfy psychological intuition that these are the best primitives even for the domain for which they were created. Some of the deficiencies may be shown by contrasting the CD primitives with some of the suggested aspects in section 5 below. What is loudly called for is a more formal and flexible representation, one which is general enough that a system can be created to learn new words and concepts in terms of these basic building concepts, without the constant aid of a programmer. An excellent book, *Conceptual Structures* by John F. Sowa, contains a more formal system (see [Sowa1984], especially in his Appendix B). The system presented there, however, while being more general than the early CD work, still does not satisfy this need: that work's "conceptual catalogue" is *not* an attempt to give a comprehensive or axiomatic set of primitives. One of our goals is to find just such a set.

It is hard to deny that in the human mind the underlying structure bears a direct functional relation to the information gathered from our senses. For example, one occasionally hears the remark that some text "brings forth an interesting mental *picture*." Separate studies have shown that reading stories bring up mental structures which behave much like one is seeing the scene before their own eyes, even when the text is describing situations with which the reader has no previous experience. The relation to the visual system is most significant, but relations to the auditory are also observable. ([Sowa1984] and [Marr1980] give many results which bear this out.)

Turning back to syntax, any ATN chart parser or similar beast, will not fully capture certain aspects of human language interaction, unless largely reworked in its means of operation. One simple example is a psychological fact that motivated Michael Dyer's program DYPAR : as we read a passage, we build up expectations as to what is coming next. We often know exactly what type of word goes in a sentence in a given location, so that we can learn new words from context. (See [Dyer1983]) Another issue involves the role of *metaphors* in language. Current systems are too stiffly formal to deal with the flexibility of words and concepts as we use them. We use words as very malleable tools for communication and mostly strive only to make sense and to understand one another, rather than to be grammatically correct. Finally, Winograd lists three further limitations to what he calls "the computational paradigm of linguistics" (once again, see [Winograd1983], this time, pages 28-30). These apply to nearly all computer NLU systems to date:

1. The many *social* aspects of language have been ignored.
2. The *evocative* aspects of language have not been dealt with.
3. The *historical* viewpoints of *language evolution* are left out.

These three limitations involve very deep questions that may never be answered via development of computer NLU models, but it is believed that the suggestions below will bring computer NLU a step closer to addressing these issues, if only because the model will more faithfully represent the methods people use to deal with language.

Most language understanding research has largely ignored the question of the language learning process until recently.[1] Chomsky pointed out that learnability is an important test for any proposed model of language competence, but then later side-stepped this issue himself somewhat by proposing a specialized *language acquisition organ* in the brain. There has been at least one major work by Wexler and Culicover [1980], searching for formal principles of *syntax acquisition* in the Chomsky school. Certainly none of the systems which we have mentioned will model the child's real-world language learning process at all. None of them even exhibit much in the way of simpler, adult language-augmentation skills, such as language acquisition from context, or from conversation, or dictionaries. This is an important lacking for at least three reasons.

First, we know that there is no single language of English. It exists as it is learned and used by native and foreign speakers, so that we will never capture it in a static set of programs and data structures. Any general-purpose system will have to be adapted to the particular technical jargon, favorite usages, the dialect, and even to some extent the *world* in which its users reside. Only a system which can easily build on its base can possibly meet such needs.

Second, we mentioned above that if it is too difficult to build the word definitions in a system, it will

---

[1] It seems odd that Schank pointed to this as an essential research goal, at least as early as 1972 - see page 629, [Schank1972] but evidently never returned to addressing it *in general* to this date. The aforementioned works have addressed the problem of learning *knowledge* from NL dialogue (as Katz and Winston also did) and of improving *story understanding*, but not of language acquisition *in general*. Other researchers have studied *syntax* acquisition, but only as a subject artificially divorced from the parallel activities of *semantic* and *pragmatic* acquisition. This author feels that many of their conclusions are in question because of this flaw. Work by Anderson [1977] and recent work by Nicholl [1987] deals with the acquisition of the semantic content of words. This work is rather promising, but has only scratched the surface to date.

be immensely difficult to implement it for any large vocabulary, hence it will be useless for many important, general-purpose tasks. A system which is adept at picking up new information would be a most welcome improvement.

Third, it is a reasonable opinion that if any machine is ever to understand language as it is understood by people, it must learn it in much the same way as people learn it. A similar idea was put forth by Weizenbaum ([1976], page 213) though with somewhat more force and generality:

> We, however, conclude that however much intelligence computers may attain, now or in the future, theirs must always be an intelligence *alien* to genuine human problems and concerns. [the italics are Weizenbaum's]

This remark is primarily claiming that there are some things machines will never do, and his reasoning to support it, simply put, was that no machine builds up a knowledge of the world in a way that resembles human growth, development, nurture, socialization, etc. I readily agree with Weizenbaum's remarks *in principle*. He has pointed out a significant defect in all of the current approaches. I am hopeful, however, that his ultimate conclusion is incorrect, since I believe such defects may be alleviated at least to some acceptable level.

Appendix A gives a summary of all of our observed limitations. These are listed together with an identification of the apparent underlying cause of each. (Some of these causes have been identified above, others have not been argued here.) What is most encouraging is that all of them can be traced in the main to weaknesses in just four areas:

1. The Underlying "Thought" Representation
2. The (Lack of an) Overall Discourse Model
3. Weaknesses in the Parsing Method
4. The Lack of Learning in the System
   (several useful kinds of learning will be identified later)

In the next section we turn away from all of this rude and negative criticism and submit some positive, constructive ideas.

§ 3. SOME GOALS FOR A NEXT GENERATION NLU SYSTEM

It would probably be inappropriate to consider (at one step) a system which had none of the many limitations described above. By considering certain basic improvements to the state of the art, however, one can imagine a system in the near future that would make significant progress in the following twelve goals. These summarize a large chunk of what has been discussed.

1) To be able to handle all of English discourse in principle, including not just simple grammatical sentences, but difficult ones. (Well, we don't really mean *all* of it, but about as much as a typical adult knows, for starters. Especially we want to handle modals, metaphors, idioms, etc.) Further,

to deal reasonably with ungrammatical but meaningful utterances, such as "How we get to Boston?" or fragments, such as "Impossible!" which are perfectly correct in an appropriate discourse, although they are not actual sentences.

2) To capture and make use of the significant linguistic regularities of the natural language. (e.g. We do not necessarily recognize *past participle*, or *embedded sentence* with grammar-school-teacher adeptness, but learnable rules with the power and generality such as those a typical adolescent might use are sought.)

3) To display predictive capability as one scans a text and to have some ability to determine lacking information from context.

4) To separate the universals of language (e.g. what kind of roles in a sentence or thought can be filled only by an animate being) from the peculiarities of any one language. (e.g. the gender of an article must agree with the gender of the noun, in most of the Romance languages.)

5) To be capable of both understanding and generating language, by closely related mechanisms.

6) To produce a deep language-independent conceptual understanding of the text it reads as a result of the understanding process. And to generate language output from such a deep-level understanding. To be capable of combining the two processes to carry on a meaningful dialogue.

7) Insofar as possible, to use methods and representations consistent with the findings of psychological and physiological research. For example, as Charniak and McDermott suggest in [Charniak1984], our representations of knowledge gained by and used in NLU will hopefully interact easily with the representations used by a vision system (they may even coincide). This goal is important if we expect later work to be able to use our NLU system in a more integrated intelligent system which will combine work from several of today's research areas.

8) To be capable of performing powerful inferences from the deep structures it works with, producing new structures. In general, this is a requirement that the deep structure be in such a formalism that inductive learning methods, deductive inference and retrieval systems, and analogical reasoning systems are able to work with it.

9) To be capable of performing high quality translation (as defined by the performance of a typical human translator in terms of on-paper competence, not other performance measures such as speed or tact, etc.)

10) To be able to paraphrase what it processes. Also to recognize at least some of the implicit information which is conveyed in passages, whether it is given as an explicit proposition or not. (e.g. When one says: "Herman was so unusually tall that he could reach a basketball rim without jumping." the main proposition is that Herman is unusually tall, and a part of that proposition is that one can gauge this in terms of a known height, that of a basketball rim. A foreigner who has

never heard of basketball however, would not fall far short - excuse the pun - of our understanding of this sentence. *In addition* he or she would probably add two facts to their knowledge: that a basketball rim is a thing and is set at a specific height in our culture, plus that said height is somewhere between five and twelve feet, in all probability. They would even stand a good chance of guessing it more accurately. Our system should be able to make at least some of this type of side inference and to store such implicit learning for future use. A danger here stems from the weakly humorous one-liner "Herman was so tall, he could reach a doorknob without jumping." This would mislead such an inferencer, but I conjecture that it would have the same effect on foreigners who did not know what our word "doorknob" meant. Such misguidances are hopefully rare enough that a robust learner gets over them using lots of separate inputs.)

11) To be *able* to learn in *all* of the following ways:
    (a) To learn the information it has been given in the discourse (see comments above),
    (b) to learn from induction and deduction, using the information in the discourse,
    (c) to learn new words, phrases, and idioms from the discourse, including new words which will imply constructing entirely new concepts (within limitations!)
    (d) to learn the structure of the language's grammar itself,
    (e) to learn multiple languages and to keep them straight for reliable communication.

12) To display a high degree of psychological validity, in that it models human thought processes and communication methods more faithfully than previous NLU systems. To make some attempt to model the other speaker in a dialogue, which will increase disambiguation abilities (cf [Grice1975], pp. 41-58).

The central thesis of this paper is that there are certain improvements which can be made in **representation, discourse modeling, parsing,** *and* **learning** *which will have a dramatic impact, easing many of the limitations discussed in the previous section and pushing forward many of the goals just enumerated in this section.*

The remaining sections deal with these four areas.

§ 4. THE UNDERLYING REPRESENTATION

It has been said that the most important task for AI in the current decade is knowledge representation. Many novel knowledge representation forms have been used for programs, a few of which have already been mentioned. A fair number of large and ambitious schemes have been developed in the AI community: Minsky's frames, Schank's Conceptual Dependency Diagrams, Quinlan's Semantic Nets, Slot-Filler notation, classical FOPC, the Horn Clauses of Prolog, extensions to FOPC such as Michalski's APC, Michalski and Winston's VPL, the Modal and Nonmonotonic Logics invented for planning systems, Winograd and Bobrow's KRL, and the knowledge form that Lenat has chosen for CYC. This last is the most ambitious to date in terms of its scope. In the last few years there have been many comparisons and debates and results proven about various of these schemes, but certainly no clean, "axiomatic" set of primitives has yet emerged. The larger question here is whether or not we can model all of human thought in *any* method.

This question has been the study of philosophers (Aristotle, Leibniz, and Wittgenstein) and has been profitably addressed in psychology (most notably Piaget, also Cohen, Rummelhart). The question of representation has even been studied via physiology. (This author has not seen any results applicable to the Artificial Intelligence community yet, except for some intriguing results in studies of vision.) Finally, the question of representation has received the attention of researchers in AI. In the study of language understanding, we need not seek representations that match neural behavior on the cellular, "hardware" level, but rather a scheme to cover human thought on the conceptual, conscious level.

The general idea of an *object-oriented language* appears to be right on the mark. By this it is not meant that *communicating modules* with *message-passing* is necessarily the correct approach, but certainly human thought does seem to focus on *objects* and *classes* which can be viewed at various levels of abstraction. These seem a satisfactory way to capture the basic *stuff* of conceptual thinking. (The aforementioned CYC project going on at Microelectronics and Computer Corporation under the direction ofDoug Lenat is building a massive heirarchical set of classes for objects.)

In such a division, the basic conceptual unit could be referred to as a *thing*, which could then be broken down into *physical* things and *mental* things). Since this all exists solely within the processor as it models itself and its world, these are *all* "mental things", strictly. (That is, in reality when the program has a concept of the "thought" class, it is "thinking about thinking", to be pedantically correct. To simplify the discussion, we will drop this first unuseful reference.) The notions of *algorithms* and of *acting processes* should be basic as well. A process could be modeled as a time-sequence of states, but it is not yet clear whether this is a good way to capture people's intuitive ideas about processes. There is an active debate now concerning several competing proposed methods for organizing temporal processes, events and relations.

Below are listed seven classes of basic concepts and conceptual relations. These provide a set of building blocks from which other concepts can be constructed.[2]

1) A world-model device for settings: x, y, z, axes and t as well.
   Abstract ordered dimensions are also available. The concept of infinity is incorporated into this.
2) All objects can be manipulated using a powerful set theory.
3) Relations can be defined and used (functions are a special case)
   Other important special cases: the FOPC functions, operators, predicates.
   Causality is a particularly useful relation.
   Equality, greater than, etc. are all necessary.
4) Mathematical structures can be constructed and manipulated. For example:
   Groups, Fields, Matrices
5) Logic is available on appropriate objects. Modus Ponens, Inheritance, and an advanced set of quantifier types and variables are available.
6) Attribute Bundles Some distinguished attributes:
   "true", "believed", "supposed", "known", etc. ...
7) The aforementioned dichotomy between mental and external things keeps us straight on issues of "real" vs. "imagined", "planned", "desired", etc. A tough question to be answered is: How many of these mental states involve distinct primitive things? anger, hate, joy, love, hope, insanity, emotionalness, confusion, drowsiness, ... John Hagueland begins to chart this difficult territory in his survey work *Artificial Intelligence: The Very Idea* ([Hagueland1986]).

The underlying representation as described above should be forced and reworked as much as necessary to make it behave much as if one were visualizing the actors, actions, and relationships in the discourse being understood. If this could be accomplished, we would immediately meet several of our goals.

One of the methods used to identify the above seven elements involved the selection of some 200 words at random from Webster's *New Twentieth Century Dictionary*. Three example definitions are given here. The first is a fairly basic, concrete term, the second is somewhat problematic, and the third represents one of the most difficult to represent. Actually, all three are easy to represent on a gross level, but the depth and abstractness increase. We will not give a representation in any suggested formalized language, since that would require a mass of explanation. But we will give definitions which show the issues to be

[2] A true purist might ask the system to first *learn* that these are useful. But a search for ultimately fundamental and universal principles of thought can quickly lead to a circularity. It can be worthwhile to a point, and philosophers have tackled the issue on occasion, but we do not wish to pursue any such thing, for fear of becoming like the dog who chases its own tail - The currently available evidence supports the notion that infants do *not* start life as such thorough *tabulae rasa*, but already have many many reflexes, schemes, and at least predispositions. We choose to give our systems such a head start. In fact, we will want to start much farther along in development, to avoid being overly ambitious. Thus the ready-made catalogue of primitives and their operations.

grappled with. Hopefully the reader can imagine casting these definitions in a more formal way using the seven types of machinery just listed. Even when formalized, our representations will not be canonical - many different ways can be devised to define these concepts within our scheme.

Example 1: The word *scythe*

scythe := NOUN, Class name for an object, 3-D, structured.

Parts: (a) Handle 1, (b) Handle 2, (c) Shaft, (d) Blade

Function: Users = Human Farmers

Use = To cut crops, e.g. wheat
See Harvest (a supergoal)

Method = 1) The scythe is held by the handles for manipulation.
2) The scythe is swung in an arc (more detail given here).
3) If the blade impacts a soft thing, it is severed (more here).

Structural Model:

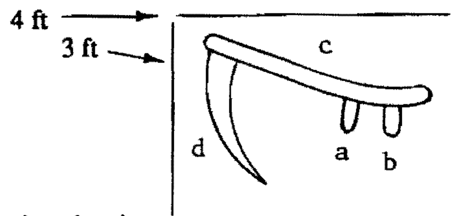

This first example is typical of a definition for a concrete noun. The main point here is that there are certain slots which are often filled for concrete nouns. There is a reasonably small set of available slots and each one has specific meaning to the system. They are not simply open in their interpretation, as would be the case in most links in a semantic network. Instead they are to be given well-defined operational meaning.[3] In most examples of semantic networks in the literature, only the "isa" link seems to enjoy this definiteness.

It seems appropriate to actually store a structural model of objects, to enable various sorts of reasoning to be done. The model would best be cast in the language of a geometric modeling system. Although not shown here, certain properties need to be attached to the geometric model, such as the fact that the shaft is made of wood and the blade of a sturdy metal. We easily take such things for granted, but they nonetheless are there in our own mental models. It is because of this added information that we have termed it a structural model rather than a geometric one, although current geometric modeling systems available have such capabilities. The actual form of this representation in the computer should match the forms the system uses for perception, just as our mental models share the forms of our perceptions.

In order to simplify, the "method" given above is very brief. In actuality, it would have to be given in considerably more detail. What is needed is a description which gives the system enough information to carry out a *scything script*. But more than a script is desired, so that there is the possibility of creating new uses for a scythe in novel situations where one is handy. Again, the formal storage of the "process of scything" should be closely related to the perceptions that would be stored if the system were "watching" harvesters scything.

[3] In the terminology of object-oriented programming, what is being suggested in part is that *methods* will be defined which operate with these *slots*.

Example 2: The word *alive*

alive :=
ADJ, Attributes include a dynamically stable biological process. (e.g. a *tree, person, bird, dog, bacteria* Not sure: a *virus*) Note that it may be necessary to distinguish that a decomposing process does not qualify, so that a dead and rotting tree, while it contains life, is not itself considered alive.

biological :=
ADJ, composed primarily of biochemical molecules (a subclass of organic molecules) in such a manner that metabolic processes are possible. Note that this definition does not give the details of what kind of composition will allow for metabolism - so in a sense our definition of biological is incomplete (which is why we still study biology!) It can only be determined that a structure can support metabolism by observing metabolism at one time in the structure.

metabolism :=
NOUN, process. "X is metabolising." A process in which X ingests matter containing stored chemical energy and breaks it down to make use of the energy in other ways, egesting the used matter. The energy is used for locomotion, growth, and maintenance of the system X. Alternatively, matter is brought in, light energy is used in photosynthesis and stored in molecules. The energy-rich molecules are again used as before.

Some other meanings which are not explored in as much detail here :

alive :=
ADJ, Attributes include a presently functioning mentality. (e.g. *God*)

alive :=
ADJ, a color pattern has many bright colors or causes lots of eye movement typically when a person looks at it..sp .2.ip "alive :=" 3 ADJ, a piece of music has a great deal of "movement". That is, the score of the music shows more jumps and notes per unit time than usual.

This example is fairly challenging. Note that in some cases our definition must remain incomplete. For example, we may not be able to teach the system all about metabolism or photosynthesis. The definition can stand even without completely defining everything as precisely as we would like. Indeed, our own concept of *alive* as humans is just such an incomplete thing. It is based in part on some examples (like those given above) and our abilities to make powerful generalizations, even when we have trouble articulating the properties of our generalizations. Our own definitions are also frequented by *black boxes* such as "photosynthesis" is in our example above. The presence of black boxes should not be cause for concern, indeed it is an intriguing facet of the system: it seems that it would not be very difficult to get such a system to display some curiosity in certain natural ways.

Example 3: The word *merciful* [4]

merciful := ADJ, attribute: history includes multiple actions of mercy.

mercy := NOUN, action (ALLOW) not requiring just punishment to be carried out.

We will not go into depth with this one, but note that it relies on several other slightly simpler, but still very complex concepts, such as justice and positions of authority in social groups of people. These in turn are built on slightly simpler concepts. For instance, justice is based on the simpler concepts of fairness and of crime and rights. It also relies in our usage on having a large body of actions which are labelled *criminal* or merely *wrong*. From these basic inputs, we inferred many rules to classify them and justify classifications as we grew up. We have a whole pyramid of concepts which build up to define *merciful* and this one is used in defining other concepts as well. It is this phenomenon which has evidently led Winograd, Bobrow, Lenat, and others to suggest that we simply build a semantic net and any concept's meaning is defined solely by its position (primarily its *connections*) in the net.

But as Piaget's work made very clear, there is an order to these things. There are basic, early concepts which have sense-related and concrete operational meanings. Our higher concepts are grounded in these. This writer feels that such a grounding is important to the success of any grand representational scheme. Otherwise, definitions are always circular and looping behavior is a constant danger. Although our three examples look much like dictionary definitions, our main point is that every word will be broken down eventually into the concept building blocks given in the list of seven categories. Even while examining my short list of 200 words, the basic seven categories have undergone major overhauls. The current list is certainly not going to be the final form. But the list has stabilized considerably. I am optimistic that a manageably concise list will succeed eventually. These few things can be given the perceptual and operational definitions, and all the rest of our concepts can be expressed in terms of these. This is a point of faith, to be sure, but it seems a well-founded faith.

Let me re-state for emphasis two important facts here. First, *this system has some things in common* with Quinlan's semantic networks, with Bobrow and Winograd's KRL and with Lenat's CYC knowledge forms. Most importantly, it shares the properties of immense extensibility and of self-definition with them. And like the CYC system, we need to build a lot of learning functionality into the system so that it can help build itself. But second, *this system has some important differences from the epistemological viewpoint:* (a) There is a harmony between this and the sensory input schemes that are expected. (b) There is a basic level that is grounded in operational and sensory knowledge. These two items will give the knowledge base a much firmer ground in the real world than a disembodied semantic net style of knowledge can achieve. (c) Certain types of actions will be defined for only certain kinds of objects - typical of an object-oriented programming environment.

---

4 The actual word which was randomly selected was "merciable", a rare form having identical meaning to the word "merciful". Mechanisms for exact synonyms are trivial to implement. We have given the more common form to avoid unhelpful distraction.

§ 5. TOWARDS DISCOURSE MODELING

We will not discuss the subject of discourse modeling except to make a couple of brief observations. Certainly any system which claims to understand a nontrivial passage will have to build a representation of the information in the passage which is built up from the information gathered from individual sentences. If the passage is a story, there will be a sequence. If the passage presents an argument, there will be a chain of inferences or causes and effects. If the passage is a detailed description, there will be a higher view to assemble from the smaller pieces of information. Also, much can be gained by knowing how the reader, the writer, and the text relate to one another. Thus it is valuable to add these elements to the overall picture. Appendix B gives an example passage of text that involves some significant issues of representation. This provides another illustration of our proposed representational scheme and involves discourse modeling.

§ 6. REGARDING THE PARSE STRATEGY AND THE LEXICON

Exception hierarchies are already in use in a number of areas of AI. They are in principle quite like the decision tree structures on which rule-based expert systems are built. They are also equivalent to a restricted form of search tree, and are of great utility in a very large class of problems. It is here proposed that each word we have learned in our experience be stored using an exception hierarchy. If the word has a meaning (such as the word "bush"), then that default meaning is stored at the top node indexed by this word. Then, if the word participates in any known larger formulas in ways that are not directly predicted by its unit meaning, these are added to the tree for each word in the formula. If the word has no meaning when standing alone (such as the word "the") then there is no default meaning in the exception hierarchy. (But there may nonetheless be properties - this may get as messy as today's ATN's have!)

For example, "the bush" is a larger formula, and is a special case of the formula "the N" where N is any noun. This formula occurs under the word "the" but not under the word "bush". The basic meaning of "bush" will be retained, but will be slightly modified using the *meaning rule* attached to the formula "the N". More interestingly, the idiom "beat around the bush" is indexed by "beat", "around", and "bush". This is because it does not preserve the most likely meanings of any of these three words. (Note that our exception hierarchies will thus form a lattice structure.) In such matching, the *largest* exception rule (deepest in the lattice) always prevails. It is never necessary for such a system to do backtracking, but it will probably prove slightly more efficient to do so in rare situations. These situations will correspond roughly to the *garden-path sentences* which have been discussed in the literature. These situations should not be viewed as a weakness in the model; this author views them rather as a selling point, for they support the notion that this system has psychological validity.

The formulas that could potentially be used in the hierarchy lattice should not be very restricted. For example, we do not propose that they be restricted to including only the usual constituent sturctures that are commonly used in linguistics, though such rules are one potent means of consolidating formulas. In addition, we will allow for formulas that involve other aspects of life that are not principally linguistic in nature. For example, there may be a rule indicating that "watch X" means something if X is a person, thing, or event, unless the thing is inert (no action at all), in which case "watch X" has a different expected meaning. For example, we have one meaning that comes to mind for the phrase "watch the performance", but a rather different meaning for "watch the shop". [5] Having noted that there are a broad range of formulas that will be needed for building the recognition of meanings, we should also note that a broad range of *meanings* should be allowed in the nodes of the hierarchy. Recall our earlier example of the node "the N", where N is any noun. We don't really want to store a meaning in this node at all, so much as a formula explaining what to do with the meaning of N. For the difference in usual discourse between "boy", "a boy", and "the boy", is really a matter of how specific we are being as to whom we are discussing. We give an example set of formulas in Appendix C. This set could have been used when interpreting the example passage in Appendix B, which would lead to the representation given there.

§ 7. LEARNING

Because of their informality and simplicity, exception hierarchies are able to be learned incrementally. A good induction routine can work on noticing regularities and turning them into higher concepts. This work could be built from some of the programs developed by the Intelligent Systems Group here at the University of Illinois. (Specifically, the programs INDUCE and CLUSTER would be suitable. See [Michalski1976] and [Stepp1986], respectively.) The hierarchy should store *frequency-of-use* information at each node which would drive the most common definitions to the top. This model could thus be used for the study of certain language acquisition processes. It seems reasonable as a psychological model, a claim which at this point is based solely on my introspective experience.

Most classical learning programs, such as those mentioned just above, can perform groupings or learn rules, if they are first given a specific set of attributes or relations to consider and then given a set of example events in the space of these attributes. In language learning, a lot of learning can take place in such a framework, but a child has no such prepared environment in which to work. The relevant attributes to observe must come either from predispositions in the child's mind, or from a very creative part of the human intellect which has not yet been explored in the machine learning research to date. It is possible that our seven-point knowledge representation could provide a restrictive enough framework in which relevant variables to observe would be able to be selected from a manageable list of candidates in spite of the rich generality of the system. Specialized heuristics for generalization and pattern recognition may be needed to make certain types of learning feasible.

§ 8. SUMMARY

We have surveyed some of the best-known works in the field of NLU by computers, and found that while they represent exciting accomplishments, we have only begun the long journey that may one day lead to programs that exhibit *in depth* understanding. From the many possible observations which could be

---

[5] This example illustrates potential garden-path behavior: "watch the shop collapse" is a reversal from the special case to the more general. This could either be handled as an *exception to the exception*, or via backtracking. Of course, this is not a very good example of garden-path behavior, since by either method, the erroneous interpretation is easily and quickly recovered from. Truly confusing sentences involve misleading matches that are more significant, such as: "Throw the man overboard a lifeline." Note here that a real investment may have been made in building a mental model by the time we have read "Throw the man overboard". →

made about the state of the art in NLU, we chose a list of limitations and future goals which seemed most appropriate as the next frontier which can be pushed - a set of goals that appear both within reach *and* appear most generally useful. Appendix A summarizes these notes. Finally, we put forth some specific ideas which could be incorporated in a system which could be built with a few man-years of research investment, which would meet a number of these goals in natural, efficient ways. It is hoped that such a system can be built, and that it would make a significant positive contribution to our understanding of the natural language process. Appendices B and C give an example to illustrate these specific implementation ideas.

It is our hope that such a system would demonstrate much greater capabilities for truly understanding natural language, and at modeling the human language capabilities (including language acquisition). Such a system holds tremendous promise for some current goals that have been set before the artificial intelligence community as a whole, such as reliable automatic translation and inference. It also holds great promise in providing the kind of human machine interfaces which have been a long-standing hope and dream of many in the field.

This model needs to be restructured when the next two words are read, so the garden-path behavior is more apparent here.

REFERENCES

Anderson, J. (1977) Induction of Augmented Transition Networks. *Cognitive Science* Volume 1, pages 125-157.

Aristotle, see for example: *The Categories*, *Posterior Analytics*,
or *On Interpretation*,
Available from the Loeb Classical Library, Harvard Univ. Press, Cambridge, Ma.

Charniak, Eugene (1984) and McDermott Drew *Introduction to Artificial Intelligence*, Addison-Wesley, Reading, Ma.

Chomsky, Noam (1957) *Syntactic Structures*, Mouton & Company, the Hague

Chomsky, Noam (1959) "On Certain Formal Properties of Grammars", *Information and Control*, 2, pp. 137-167

Chomsky, Noam (1965) *Aspects of the Theory of Syntax*, MIT Press, Cambridge, Ma.

Dyer, Michael (1983) *In-Depth Understanding*, MIT Press, Cambridge, Ma.

Gentner, Dedre and Gentner, Donald (1983) "Mental Models of Electricity", pp. 99-129 in *Mental Models*, Laurence Erlbaum Associates, Inc. Hillsdale, NJ

Grice, H. Paul (1975) "Logic and Conversation", *Syntax and Semantics 3: Speech Acts*, edited by Cole, Peter and Morgan, Jerry, Academic Press, New York

Halliday M.A.K. (1961) "Categories of the Theory of Grammar", *Word*, 17, pp. 241-292

Katz, Boris (1982) and Winston, Patrick *Parsing and Generating English Using Commutative Transformations*, M.I.T. Artificial Intelligence Laboratory A.I. Memo No. 677

Marr, David (1980) *Vision*, M.I.T. Press, Cambridge, Ma.

Michalski, Ryszard (1976) Chapter 4, from *Machine Learning (Volume 1)*, edited by Michalski, Carbonell, and Mitchell, Tioga Publishing, Palo Alto, Ca.

Nicholl, Sheldon (1987) "Language Acquisition by a Computer Program based on First Order Logic. M.S. Thesis, University of Illinois at Urbana-Champaign

Schank, Roger (1972) "Conceptual Dependency: A Theory of Natural Language Understanding", *Cognitive Psychology*, 3 pp. 553-554

Schank, Roger (1973) and Colby, Kenneth, editors *Computer Models of Thought and Language*, W.H. Freeman, San Francisco

Schank, Roger (1982) "Reminding and Memory Organization: An Introduction to MOPS." In *Strategies for Natural Language Processing*, Laurence Erlbaum, Hillsdale, NJ

Schank, Roger (1984) *The Cognitive Computer*, Addison-Wesley, Reading, Ma.

Sowa, John (1984) *Conceptual Structures*, Addison-Wesley, Reading, Ma.

Stepp, Robert (1986) and Michalski, Ryszard "Conceptual Clustering: Inventing Goal-Oriented Classifications of Structured Objects", *Machine Learning II*, Morgan-Kaufmann

Weizenbaum, Joseph (1976) *Computer Power and Human Reason*, W.H. Freeman, San Francisco, Ca.

Winograd, Terry (1971) *Procedures as a Representation for Data in a Computer Program for Understanding Natural Language*, MAC TR-84, MIT Art. Intel. Laboratory, Ph.D. Thesis

Winograd, Terry (1983) *Language as a Cognitive Process, Volume I: Syntax*, Addison-Wesley, Reading, Ma.

Wittgenstein, Ludwig (1953) *Philosophical Investigations*, Basil Blackwell & Mott Ltd., Oxford

Wittgenstein, Ludwig (1974) *Philosophical Grammar*, Univ. of Calif. Press, Berkeley, Ca.

APPENDIX A

Summary of Limitations of Current NLU Systems, Labeled by Cause

Key: **PS** = Parsing Strategy **UR** = Underlying Representation
DM = Discourse Model **LE** = Learning

Limitations of Present Computational Linguistics Work	
1. The Syntax versus Semantics Debate :	
a. Failure to Predict Human Parse Times	**PS**
b. Interpretation While Scanning (Expectations)	**PS**
c. Sparse Information in Familiar Situations	**DM**
d. The Inference Problem	**UR, DM**
e. Paraphrasing	**UR**
f. Translation	**UR**
g. Ambiguities	**PS, UR**
h. Pronoun Reference	**UR, DM**
2. Hypothetical Statements	**DM**
3. Counterfactual Statements	**UR**
4. Limited Declarative Information: Don't Represent Time, Causality, Modality, etc.	**UR**
5. Verbal Acts	**DM, LE**
6. Model of the Other Conversant	**DM**
7. Model of the Conversation Process	**DM**
Limitations of Current CD-Based NLU Systems	
1. Syntax versus Semantics (revisited)	
a. Do Not Capture "Significant Linguistic Generalities"	**PS**
b. Inability to Handle Many Sentences	**PS, DM**
c. Difficulty Extending	**PS, UR, LE**
2. "CD was Never Meant to be General"	**UR**
3. CD is Not Even A Good Basis From Which to Extend	**UR**
4. CD Will Not Ultimately Provide a Psychologically Valid Model	**UR, DM**
5. The Language Understanding Problem Has Been Shifted, But It Has Not At All Been Solved	
a. Consider the Policeman in a Restaurant	**PS**
b. Ambiguity Problems Remain, Including Pronoun Reference	**PS, DM**
6. Probably No Hope On the Horizon for Abstract Conversation	**UR, DM**

Limitations Common to All Current Systems	
1. Cannot Yet and May Have Real Difficulty Ever Handling:	
a. Large Vocabulary	**UR, LE**
b. Complex Discourses	**UR, DM**
2. Psychological Validity	
a. Mental Pictures	**UR, PS, DM**
b. Relating to Other Areas:	
i. Vision	**UR**
ii. Hearing	**UR**
iii. Logic / Inference / Mathematics	**DM, LE**
c. Studies of Language Performance Errors	
i. Nonmeaningful Discourse Pauses, Filling Errors	**DM, PS**
ii. Garden Path Sentences	**PS**
3. Metaphor and Abstract Thought In General	**UR**
4. LEARNING!	
Note: language learned for *meaning*, not just formal syntax	
a. Learning Knowledge from the Discourse	**LE, UR, DM**
b. Learning Words	**LE, UR**
c. Learning Language Syntactic Rules	**LE, PS**
d. Learning Idioms	**LE, PS**
e. Learning Strategies for Meaning	**LE, UR, DM**
f. Learning Strategies for Understanding (Parsing)	**LE, PS, UR, DM**
g. Learning Multiple Languages for Translation	**UR, PS**
h. Learning Evocative Aspects (Artistic, Stylistic)	**UR, DM**
5. Winograd Listed Limitations: (re: All Current NLU Systems)	
a. Social Aspects of Language Have Been Ignored (Functional Grammars begin to deal with this)	**UR, DM**
b. The Evocative Aspects of Language Have Not Been Dealt With	**UR, DM**
c. The History and Evolution of Languages Are Not Dealt With	**PS, UR, DM**

If a machine is ever to understand language as it is understood by people, it must learn the language in much the same way as people learn it.

APPENDIX B

Knowledge Representation for a Complete Discourse

The following short passage gives an illustration of our proposed scheme. After the passage is a list of concepts showing the manner in which the final *thought representation* is constructed. Then the final *thought representation* is given. Although it is diagrammatic in nature, this is just for presentation. The internal representation is more mundane, though directly translatable from and into such a picture. Note that our passage involves several aspects of common human discourse that are beyond the scope of prior systems (e.g. time sequencing, emotional content, abstract concepts).

> *So Solomon observed the feast at that time, and all Israel with him, a great assembly from the entrance of Hamath to the brook of Egypt, before the Lord our God, for seven days and seven more days, even fourteen days. On the eighth day he sent the people away and they blessed the king. Then they went to their tents joyful and glad of heart for all the goodness that the Lord had shown to David His servant and to Israel His people.*

From this, the following concepts were produced. They are listed in the order produced while reading and understanding the passage. Note that this table does not include any "false parses"; concepts which may have been built but were later determined to be an incorrect interpretation of the text are not listed.

1	human, specific	9.1	set of humans, specific $= (5)$
2	seq of events, $\neg$ specific	10	physical structure, $\neg$ specific
2.1	seq of events, specific	11	human, specific
3	seq of events, specific	10.1	physical struct, specific
4	time, $\neg$ specific	12	physical object, $\neg$ specific
4.1	time, specific	12.1	physical object, specific
3.1	seq of events, specific	13	nation, specific
5	set of humans, specific (related to (1))	12.2	physical object, specific
6	human, specific $= (1)$	14	2-D region, specific
7	set of humans, specific $= (1) \cup (5)$	9.2	physically localized (specific loc) set of humans, specific
8	seq of events, specific $(3) \subseteq (8)$		
9	set of humans, $\neg$ specific (physically localized, but location $\neg$ specific)	8.1	sequence of events, specific (setting and actors known (9.2))

--- (table continued) ---

--- (continued from previous page) ---

16	diety, specific = (15)	30	action, specific
15	diety, specific	31	mental state
16.1	diety, specific = (15)	32	body part
8.2	seq of events, specific = (8.1)	33	mental state
17	time interval, $\neg$ specific	34	type of behavior, abstract
17.1	time interval, $\neg$ specific	34.1	seq of behavior, abstractly characterized
18	time interval, $\neg$ specific	35	diety, specific = (15)
18.1	time interval, $\neg$ specific	36	human, specific = (related to (1), (5))
19	time interval, $\neg$ specific = (17.1) $\cup$ (18.1)	37	human, $\neg$ specific (role characterized)
20	time interval, $\neg$ specific	37.1	human, specific
20.1	time interval, $\neg$ specific	38	relation
8.3	seq of events, time duration specified = (8)	39	set of humans, specific = (5)
21	time interval, $\neg$ specific	40	set of humans, $\neg$ specific
21.1	time interval, $\neg$ specific	40.1	set of humans, specific
22	human, specific = (1)	41	relation
23	set of humans, specific = (5)	42	relation = (38) $\cup$ (41)
24	action, specific (causal links to other actions)	34.2	seq of behavior, abstractly characterized
25	set of humans, specific = (5)	43	mental state with a cognitive object
26	human, specific = (1)	44	mental state = (31) & (43)
27	action, specific	45	action, specific = (30)
28	set of humans, specific = (5)	46	seq of events, specific, time $\neg$ specific = (24), (27), (45)
29	physical objects, $\neg$ specific	46.1	seq of events, specific, time specific
29.1	physical objects, specific	47	sequence = (8.3), (46.1)

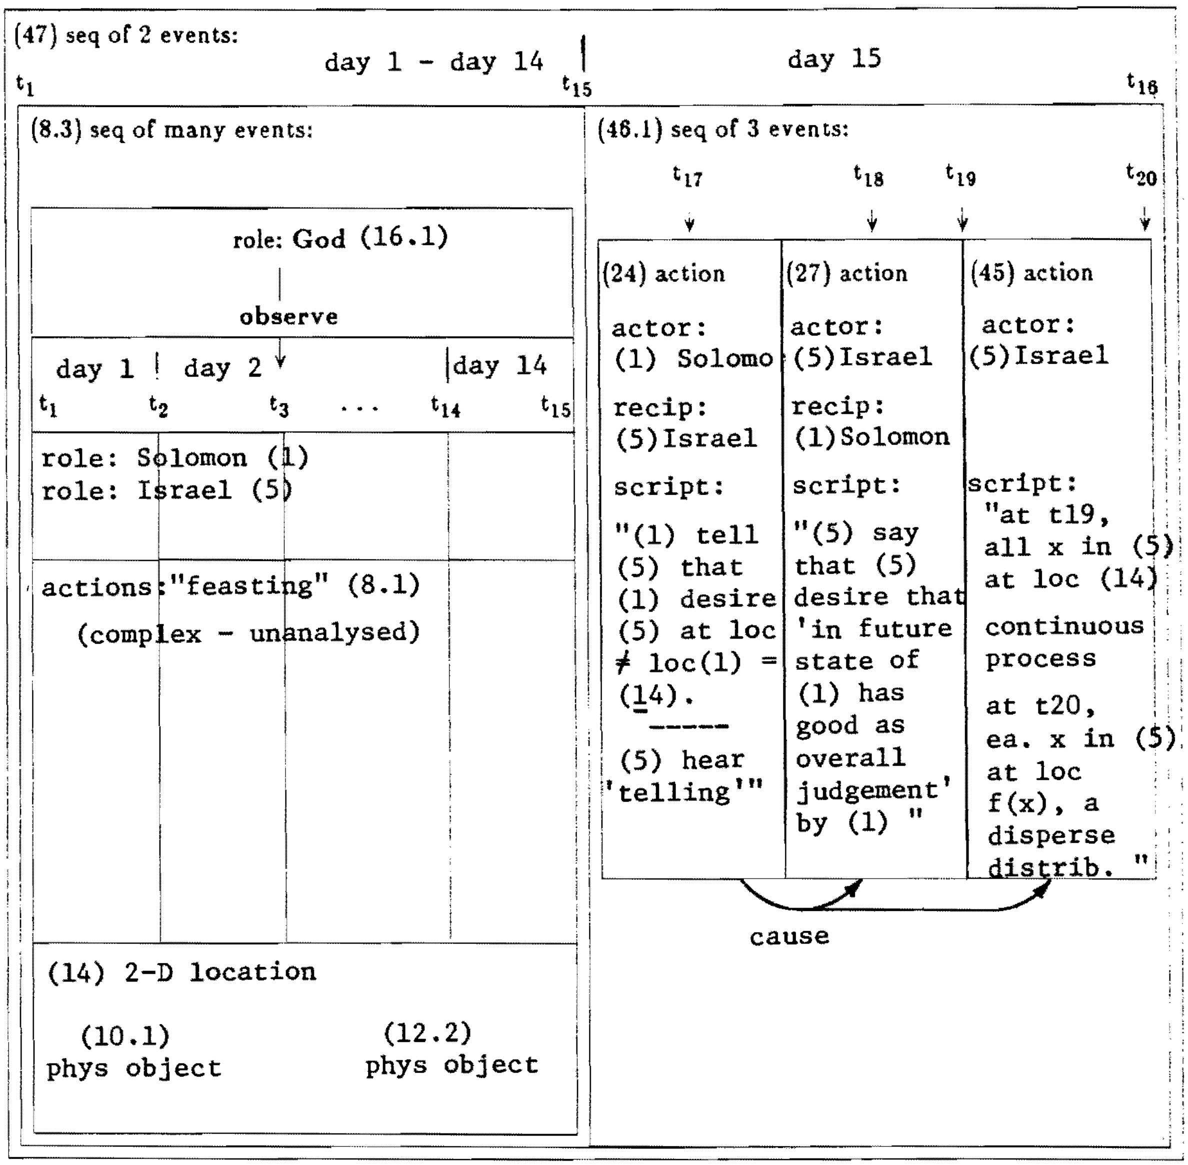

Figure B-1 : A Knowledge Representation for the Passage

Note that on the right side, we are just getting to the "fun" part as we begin to expand the three scripts which have only been very informally specified at a cursory level. Any serious treatment would have to be able to expand these on demand to a considerably finer level. But at first reading, there should be control mechanisms in place to limit the expansion to keep the tradeoffs between detail and time expenditure in balance.

APPENDIX C

A Sample Lexicon for the Discourse of Appendix B

Note that the following is only semi-formal. For brevity, many of the definitions are given in a human-readable rather than machine-usable fashion. The intent is to give the reader a feeling for the information that will need to be put in the lexicon (in a careful, painfully formal way) and its general organization. Note also that in our lexicon, many meanings of words are skipped. For example, we know many different uses for the word "entrance" and any responsible lexicon-builder will want to eventually include more than we have here. We have restricted the size to suit the immediate needs of the passage we are analyzing. It is important to point out that additional meanings are easily added. No matter how ambitious we are in our first attempt at defining a word and its uses, we will always need to make additions and alterations as our knowledge grows. In our lexicon, this is quite straightforward.

a := nil

- **a object class** := 1 (**object**) [not specific (don't seek referent)]
    - f: **object** [¬quant] → **object** [quant]

all := nil

- all **set** := **set** [high precision of inclusion]
    - f: **set** → **set**

and := nil

- **object$_1$** and **object$_2$** := { **object$_1$** , **object$_2$** }
    - f : (**object** , **object**) → **set**
- **object** and **set** := { **object** } ∪ **set**
    - f : (**object** , **set**) → **set**
- **set** and **object** := { **object** } ∪ **set**
    - f : (**set** , **object**) → **set**
- **set$_1$** and **set$_2$** := [**set$_1$** ∪ **set$_2$**]
    - f : (**set** , **set**) → **set**
- **state$_1$** and **state$_2$** := **state$_1$** & **state$_2$**
    - f : (**state** , **state**) → **state**

assembly := **set** {H_1, H_2, ... , H_0 [each H_i : human (default assumption)]
    [default n ≈ 100, n ≥ 2 required for validity, n ≥ 10 for appropriateness
    setting: L such that vol(L) ≈ 100m by 100m by 20m (default assumption)]

away := nil

- away from **X** := aspect **X** ⟶
- away to **X** := aspect □ ⟶ **X**
- **X** send **Y** away := **process$_1$** : t_4 : comm(**X**, **Y**, desire(**X**, **process$_2$**))

    process$_2$: t_1: **X Y**
    t_2: **X**   **Y**
    t_3: **X**   **Y**

at := nil

- [**event**,**object**] at **location** := loc([**event**,**object**]) = **location**

• [**event**,**object$_1$**] at **object$_2$** := loc([**event**,**object$_1$**]) = loc(**object$_2$**)

before := nil

• **event$_1$** before **event$_2$** := **seq**: t_1, **event$_1$**, t_2, **event$_2$**, t_3

• **event** before **object** := loc(**event**) = in-front-of-loc(**object**)

• **event$_1$ occur** before **event$_2$** := **seq**: t_1, **event$_1$**, t_2, **event$_2$**, t_3

bless := nil

• **person$_1$** bless **person$_2$** := **event**: say(**person$_1$**,desire**person$_1$**,state(**person$_2$**,S)
where X is judged to be "good" by narrator.

brook := **object class**

• the brook of Egypt := **object** [specific]

David := **human label** [specific]

day := **time interval class** [not specific] aligned

Egypt := **nation label** [specific]

eighth := nil

• the eighth **X** := seek countable indexed set of **X**, with ≥ 8 members.
select item 8 from the set.

entrance := • **object** part of enclosure-object through which other objects may pass.
• **event class**

• the entrance of Hamath := **object specific**

• **person** make an entrance := **event** person entered location of narrator

even := relator, =

feast := **complex event class**

t_1 person$_1$ eat t_2

t_1 person$_2$ eat t_2

. . .

t_1 person eat t_2

for := nil

• **event** for **time interval** := occurred(**event**, **time interval**)

• **emotional state** for **event** := cause(**event**,**emotional state**)

fourteen := number: 14

from := nil

• from **location$_1$** to **location$_2$** := **region** space interval [**location$_1$**, **location$_2$**]

glad := mental state = "happy"

• glad of heart := glad

God := **diety** [specific]

goodness := **event class attribute**

great := nil

• great **object class** := subclass of **object class**, significantly larger than average for that class

had := nil

• had **event** := occurred(**event**,t_1) t_1 < **time of context event**

Hamath := **human label** [specific]

he := **human** [non specific, male] invoke anaphoric resolution

heart := complex animal body organ.

• glad of heart := glad

him := nil

• **relator** him := role player in **relation** = **human** [non specific, male] invoke anaphoric resolution

his := **nil**

• **his object** := possess(**human** [non specific, male] invoke anaphoric resolution, **object**)

• **object his relator** := **object**, relation(**human** [non specific, male] invoke anaphoric resolution, **object**)

Israel := **nation label** [specific]

joyful := **nil**

• **event joyful** := emotional state(**participant(s) of event**, happy) at time-of-event

• **person be** joyful := emotional state(**person**, happy) at time-of-discourse

king := • **office** class

  • **person** [holder of office]

Lord := **diety** [specific]

more := nil

• more **object** := **object** [defeat unification with prior **object**]

observed := nil

• **person** observed **event** := observe(**person**, **event**) [primitive function]

• **person** observed **holiday** := participate(**person**, activities of(**holiday**))

of := nil

• **object** of **person** := **object** [determiner: associated with **person** (default: possession) (2nd choice: named after)]

• glad of heart := glad

on := nil

• on **time event** := occurred(**event**, **time**) [primitive]

• on **object** := part-of(**object**, "on") [primitive]

our := **nil**

• **our object** := **object** [determiner: associated with **person** (default: possession)]

people := **set** { person_1, person_2, ... , person }

  [no default n, $n \geq 2$ required for validity, $n \geq 10$ for appropriateness

send := **nil**, VERB

• **X** send **Y** to **Z** := process:

sent := **nil**, VERB: past tense of send

servant := expand to: **X** is servant of **Y** = **X** does acts caused by **Y**'s will.

seven := **number**: 7

shown := nil

• $\textbf{person}_1$ had shown **attribute** to $\textbf{person}_2$:=

  relation($\textbf{person}_1$, $\textbf{person}_2$) [attribute: **attribute**]

so := **nil** (Let's leave it at that, for this passage!)

Solomon := **human label** [specific]

tent := **complex** built object

that := **nil**

• that **object class** := specific member of **object class** [fire anaphoric resolution]

• **object** that **proposition** := specification of **object** as matching **proposition**

the := nil

- the **object class** := specific member of **object class** [fire anaphoric resolution]

their := nil

- their **object** := possess(**set of human** [non specific] invoke anaphoric resolution, **object**)
- **object** their **relator** := **object**, relation(**set of human** [non specific] invoke anaphoric resolution, **object**)

then := nil

- **$\mathbf{event_1}$** then **$\mathbf{event_2}$** := **seq**: t_1, **$\mathbf{event_1}$**, t_2, **$\mathbf{event_2}$**, t_3

they := **set of human** [non specific] invoke anaphoric resolution

time := primitive dimension

to := nil

- one entry points to the same node as found under "from ... to ..."
- another entry points to the same node as found under "... went to ..."

went := nil

- **person** went to **location** := **process** :

  t_1: **person** **location**
  t_2: **person** **location**
  t_3: **person location**
  loc(narrator) ≠ **location**

with := nil

- **$\mathbf{person_1}$** with **$\mathbf{person_2}$** := {**$\mathbf{person_1}$**, **$\mathbf{person_2}$**}

, := nil, PUNCT clause separator

. := nil, PUNCT sentence terminator (declarative)