```
*(SPRINT (TRY ROB-SUM))
********************* INPUT *******************************
((SUM ((LEAVE (AGT JOHN)
          (SOURCE LIQUOR-STORE)
          (SC (ESCAPE830 (AGT JOHN)
                  (SOURCE LIQUOR-STORE)
                  (SC (GETAWAY833 (AGT JOHN)
                          (SOURCE LIQUOR-STORE))))
               ))
       (GET1 (AGT JOHN)
          (OBJ GUN)
          (SC (GET2#759 (AGT JOHN)
                  (OBJ GUN)
                  (CONSQ (HAVE2#769 (AGT JOHN) (OBJ GUN))))))
       (ROB764 (INSTR GUN)
          (OBJ MONEY)
          (AE OWNER)
          (AGT JOHN)
          (LOC LIQUOR-STORE)
          (ANTE (TRAVEL778 (AGT JOHN)
                    (DEST LIQUOR-STORE)
                    (SUBSEQ (MOVE1#775 (OBJ JOHN)
                                (DEST LIQUOR-STORE)
                                (SC (WALK1 (AGT JOHN)
                                       (DEST LIQUOR-STORE
                                                       )))))))
          (PREC (WANT1 (AGT JOHN)
                    (OBJ MONEY)
                    (CONTINUATION (WANT2 (AGT JOHN) (OBJ MONEY)))))
          (SUBSEQ (SAY819 (TO OWNER)
                       (AGT JOHN)
                       (THM WANT2)
                       (SC (TELL1 (TO OWNER)
                                  (AGT JOHN) (THM WANT2)))))
          (PREC (HAVE2#769 (AGT JOHN) (OBJ GUN)))
          (SUBSEQ (GIVE (REC JOHN) (OBJ MONEY) (AGT OWNER)))
          (SEQ (GETAWAY833 (AGT JOHN) (SOURCE LIQUOR-STORE)))))
************************ OUTPUT *****************************
       ((ROB764 (INSTR GUN)
                (OBJ MONEY)
                (AE OWNER)
                (AGT JOHN)
                (LOC LIQUOR-STORE)))))
```

Figure 5-23: "Robbing a Liquor Store" - Summary

```
*(SPRINT (TRY PIG-SUM))
((SUM
************************* INPUT *****************************
  ((GATHER1 (AGT PIG)
            (OBJ LAUNDRY)
            (CONSQ (HAVE2#96 (AGT PIG) (OBJ LAUNDRY))))
   (APPEAR26 (PERCEPTOF PRINCE)
             (OBJ PIG)
             (LOC WATER)
             (COORD (SEE23 (THM PIG)
                           (PERCEPTOF PRINCE)
                           (LOC WATER)
                           (SC (SPY1 (THM PIG)
                                  (PERCEPTOF PRINCE) (LOC WATER)))
                           (CONTINUATION (WATCH1
                                            (THM (AND SOAK1 SCOUR1))
                                            (PERCEPTOF PRINCE)
                                            (LOC WATER)))
                           (CONTINUATION (WATCH2 (THM HANG1)
                                            (PERCEPTOF PRINCE)
                                            (LOC WATER)))))
             (PREC (COME40 (AGT PIG) (DEST WATER))))
   (CLEAN83
    (AGT PIG)
    (LOC WATER)
    (OBJ CLOTHES)
    (ANTE (MOVE2#88 (AGT PIG)
             (OBJ LAUNDRY)
             (DEST WATER)
             (SC (CARRY1 (AGT PIG)
                    (OBJ LAUNDRY)
                    (DEST WATER)
                    (COORD (TRAVEL17
                       (TWD STREAM)
                       (AGT PIG)
                       (DEST WATER)
                       (SUBSEQ (MOVE1#20
                          (DEST WATER)
                          (OBJ PIG)
                          (TWD STREAM)
                          (SC (TROT1 (EST WATER)
                                (AGT PIG)
                                (TWD STREAM)))))
                               (SC (COME40 (AGT PIG)
                                      (DEST WATER)))))))
           ))
```

Figure 5-24: "The Tale of the Pig" - Summary

```
    (SEQ (MOVE2#103 (AGT PIG)
            (OBJ LAUNDRY)
            (DEST HOME)
            (SOURCE WATER)
            (SC (CARRY166
                   (AGT PIG)
                   (OBJ LAUNDRY)
                   (DEST HOME)
                   (SOURCE WATER)
                   (COORD (TRAVEL140 (AGT PIG)
                      (SOURCE WATER)
                      (DEST HOME)
                      (SUBSEQ (MOVE1#137
                         (OBJ PIG)
                         (SOURCE WATER)
                         (DEST HOME)
                         (SC (TROT2 (DEST HOME)
                                    (SOURCE WATER)
                                    (AGT PIG)))))))))
                    (ANTE (HAVE2#96 (AGT PIG) (OBJ LAUNDRY)))))
    (SUBSEQ (WASH80 (AGT PIG)
                    (OBJ LAUNDRY)
                    (LOC WATER)
                    (SUBSEQ (SOAK1 (AGT PIG)
                                   (OBJ LAUNDRY) (LOC WATER)))
                    (SUBSEQ (SCOUR1 (AGT PIG)
                            (OBJ LAUNDRY) (LOC WATER)))))
    (SUBSEQ (DRY94 (AGT PIG)
                   (LOC WATER)
                   (OBJ CLOTHES)
                   (SC (HANG1 (IN SUN) (LOC WATER)
                              (OBJ CLOTHES)))))))
**************** SUMMARY *********************
  ((APPEAR26 (PERCEPTOF PRINCE) (OBJ PIG) (LOC WATER))
   (CLEAN83 (AGT PIG) (LOC WATER) (OBJ CLOTHES))))
```

Figure 5-24 concluded

Researcher	Summarization Base	Programmed Examples	Method
Rumelhart	Story Trees	No	By level of tree + rewrite rules.
Kintsch & Van Dijk	Macro-structures	No	Macro-structure rules.
Correira & Simmons	Story Trees Schema/Narrative Trees	Yes	By level of tree.
DeJong	Sketch Scripts	Yes	Extract from text by top-down script application.
Lehnert	Plot Units	No	Analysis of Plot Unit Graph.
Alterman	Concept Trees	Yes	Single/Multiple Concept Summarization Rules

Table 5-2: Theories of Summarization

subtrees according to a set of semantic summarization rules. Simmons & Correira work with a single tree which represents a combination of the syntactic and semantic structure of the story. Any level of their tree represents a summarization of the story. No extra rewrite rules are required.

There are a couple major differences between the work of Rumelhart and Simmons & Correira, and SUM. Instead of working with a structure which represents the story as a whole, SUM works with a series of interconnected structures (concept trees) that represent the event concepts used in the text. Where the others summarize a story according to its overall structure, SUM summarizes a story by summarizing the event concepts it uses.

Kintsch & van Dijk [Kintsch and van Dijk 75a, Kintsch and van Dijk 75b, van Dijk 76, van Dijk 77] summarize text by using a set of macro-structure rules. An example of a macro-structure rule is generalization; concepts are generalized by abstracting them. So "John is hitch-hiking" wold be generalized to "John is travelling". Another rule is deletion, which abstracts accidental properties. So "Mary was playing with the red ball" could be progressively summarized as "Mary is playing with the ball" and then just "Mary is playing". A third example of a macro-structure rule is construction, which replaces the components of a concept by the concept. So "John laid the foundation. He built the walls ..." would be summarized "John built a house".

One interesting difference between SUM's summarization rules and Kintsch & van Dijk's macro-structure rules has to do with how they would summarize two event concepts if one was a subclass of the other. Citing psychological evidence Kintsch and van Dijk choose the class concept because it represents a more general concept. SUM takes the other tack, choosing the subclass concept because it tells the reader more about what occurred in the story. The first technique results in a broad overview of the text, the second a short, more precise, abstract-like, description of the events which occur in the text.

DeJong produces representations of text by applying in a top-down fashion sketchy scripts (e.g. accidents, terrorist acts ...). His system extracts from wire-service newspaper stories just enough facts to fill in the arguments of a sketchy script. Because the stories that DeJong works with are so stereotyped he can summarize text by using a set of fill-in-the-blank type summarization statements attached to each sketchy script. The text that SUM works with is much less stereotyped and therefore a more general summarization technique is required.

Lehnert [Lehnert 81a, Lehnert 81b] has sketched out a scheme for summarizing text based on plot units. Plot units are an attempt to represent affect-state patterns. Lehnert identifies a number of basic plot units (e.g. motivation, success, hidden blessing). Associated with each basic plot is a canned program for summarization, which she calls a generational frame. Narrative text, in her proposed system, is represented by interconnected plot units. Summaries are generated by first identifying pivotal plot units (i.e. units which are maximally connected), then by generating a base-line summary from the generational frame associated with that particular plot unit, and finally by augmenting the summary with information from the plot units related to the pivotal unit.

The summarization technique she outlines is by far the most complicated. Her research program is still in the early stages so it is difficult to evaluate. Her summaries will be limited to narrative text. The same limitation applies to SUM, but SUM's could be applied to a wider range of text by extending the dictionary beyond event/state concepts. The types of summary produced by SUM differs from the types of summaries suggested by Lehnert; one bases its summaries on an analysis of the event/state concepts, the other on an affect analysis.

Chapter 6

Conclusions

6.1 Summary

NEXUS represents a theory of event/state conceptual coherence in text. The theory is: representations of narrative text can be generated by a process of matching text against a dictionary of concepts, which are related by a small set of relation-types, and using the organization of the concepts in the dictionary to organize the instances of event/state concepts which appear in the text. The dictionary was compiled by an analysis of ten folktales. The theory was tested by a program which produced coherence representations for eight samples of text. The utility of the scheme was validated by performing experiments in question answering and summarizing.

There are seven event/state coherence relations that NEXUS uses. They are: class/subclass, sequence/subsequence, coordinate, antecedent, precedent, consequent, and sequel. Class/subclass represents an inheritance relationship between concepts. Sequence/subsequence and coordinate are two kinds of part relationships: a subsequence of an event occurs for a subinterval of the event, and a coordinate occurs for the entire interval. The other four relations are temporal (see figure 6-1) and can be divided into before and after categories, and then further subdivided into necessary and plausible connections.

```
                  Temporal Relations
                    /               \
                   /                 \
               Before                 After
                /   \                 /  \
               /     \               /    \
              /       \             /      \
      antecedent consequent  consequent sequel
```

Figure 6-1: The Temporal Relations

NEXUS′ relation scheme can be used to hierarchically organize a wide range of textual phenomena. A number of interesting characterizations resulted from comparing NEXUS′ representations to other representation schemes.

When NEXUS was applied to a script example, the subparts of the script were related by subsequence arcs, and the entry and exit conditions by temporal relations.

Plans were represented by one of two techniques. Plans which were not concept oriented were represented in NEXUS by chains of temporal relations. Plans which were concept oriented (e.g. robbing) were represented in the same manner as scripts; the difference is that the default values for scripts are tightly constrained, while the default values for plan concepts are not.

When NEXUS was applied to a story tree example, "The Margie Story", a number of important features of the scheme became apparent. First, NEXUS was able to represent phenomena which neither scripts nor plans could account for; using the temporal relations it was able to connect the ′winds catching the balloon′ to the ′balloon bursting′. Furthermore temporal relations could be used to connect Margie′s reaction to the balloon bursting to the other events of the story. Finally where a story tree approach accounted for the story by constructing a tree of its overall structure, a NEXUS-type analysis produced several interconnected concept trees which represented the event/state concepts used in the construction of the text.

When NEXUS was applied to a schema/narrative tree example, "The Black and Yellow V2 Rocket", it was shown NEXUS could account for the events of the story but not the descriptive passages. This emphasized the fact that NEXUS was a theory of event/state coherence.

Finally, the comparison to Hobbs′ coherence representation explicated the difference between rhetorical coherence and concept oriented coherence. Rhetorical coherence introduces meta-relations which describe the discourse relationships between sentences. NEXUS′ concept oriented coherence captures the implicit relationships among event descriptions in the text.

The theory of coherence was tested. First, ten folktales were analyzed and used as a basis for compiling a dictionary of event/state concepts. Attached to each concept in the dictionary was a list of default values for its case arguments, and attached to each relationship a list of constraints. An interpreter, TRACE, was written which matched case notated versions of input text against the dictionary of event/state concepts and produced as an output interconnected concept trees. Both the constraints attached to each relationship in the dictionary and the default values for case arguments of concepts were used to control the representation building process.

TRACE was successfully applied to eight samples of text. Three of the samples come from the literature; these include stories which had been accounted for by either scripts, plans or story trees. The other five samples came from the folktales that had been used to compile the dictionary.

An analysis of the results provided some interesting discussion. First, it was discovered that although some paths of arc length six or seven were need to connect events in the text, the average path length between a new event description and the concept trees, as they had so far been derived, was two. Secondly, it was pointed out that in a number of cases TRACE had to use the constraints to reject several bad interpretations before it could find one that was acceptable. Thirdly, remarked upon was the fact that in the process of constructing the concept trees TRACE had resolved ellipsis and phoric references. Lastly, it was noted that TRACE had produced some novel interpretations; that is, events which had not been connected by a hand analysis of the text were correctly connected by TRACE using the dictionary of concepts.

The utility of the interconnected concept trees produced by TRACE was demonstrated by showing their usage in answering questions and summarizing text. Questions were answered by associating with each question type a list of coherence patterns that QUEST used as a basis for generating potential answers (The proviso was that the answers QUEST produced were subject to further semantic testing.). For example to answer enhancement questions QUEST descended subclass, subsequence, and coordinate relations. To answer goal questions it either ascended subclass, subsequence or coordinate arcs, or it moved forward across temporal relations.

SUM tested both the coherence relations and the hierarchical organization of the concept trees. Individual concepts were summarized by selecting, after promotion, the top node in the associated concept tree. For multiple concepts, SUM used heuristics, based on the coherence relations, to delete some of the concept trees. Two heuristics discussed were the deletion of identity trees and preparatory trees.

To summarize: the findings of this research are seven event/state coherence relations which provide a uniform approach to representing a wide variety of textual phenomena. The experiments with TRACE demonstrated that the relations could be used to organize a dictionary of concepts, which, in turn, could be used as a basis for representing instances of event/state concepts as they appear in text. Experiments in question answering and summarizing demonstrated the utility of the structure of these representations for developing simple and elegant natural language processing heuristics.

6.2 Future Directions

6.2.1 The Dictionary

There are a number of ways in which the dictionary can be extended. First by expanding the dictionary's content by analyzing, and computationally accounting for, more text. Since the dictionary is, in some sense, an evolving organism, additions to the dictionary would also, undoubtedly, refine and improve some of the current relationships between concepts.

For that matter the breadth of the dictionary could be increased. Currently the dictionary only accounts for event/state concepts. One obvious extension would be to incorporated <u>thing</u> concepts. From a NEXUS perspective this could be accomplished in two stages. First, a dictionary of thing concepts could be derived by essentially repeating the experiments on event/state concepts over a different domain of conceptual coherence. Again, text could be analyzed to derive coherence relations, and a TRACE-like procedure could be written which finds the implicit relationships between 'thing' concepts used in the text. The second stage, and the hardest part, would be trying to integrate the two dictionaries.

6.2.2 TRACE

Currently TRACE uses a smallest frontier breadth-first search across the dictionary to connect event descriptions. With the arc selection restrictions employed the size of TRACE's search space is reduced, but further reductions would help. One direction of research could be working out methods for controlling the visibility of the dictionary using the context provided by the text. Here a dictionary of 'thing' concepts would come in handy; the two dictionaries could supply contexts for each other.

6.2.3 QUEST and SUM

A number of additions could be made to QUEST. No attempt has been made to address the issue of identifying questions. For the purposes of this study it was reasonable to put the questions in a canonical form, but the concomitant problems of question typologies and question identification represent difficult research problems.

The categories of questions that QUEST retrieves answers for could be expanded. The current version of QUEST emphasized questions about relationships between events which are described in the text. That leaves out a number of categories of questions. For example, concept completion questions would represent a way to measure NEXUS' reference and ellipsis resolving capabilities.

Also, no attempt has been made to answer questions which go beyond the representation TRACE produces of the text. Ideally QUEST could retrieve answers to questions in the context of the dictionary; many questions require interfacing between the dictionary's general storehouse of knowledge and the TRACE-produced concept trees. Working out the details of an intelligent interface between the two sources of information available to QUEST is another interesting research problem.

The number of spin-off research problems from SUM are fewer than those of QUEST. One natural follow-up on SUM would be to add more concept deletion heuristics. Another interesting problem would be to integrate the concept trees for an entire story into a story tree, and then investigate summarization heuristics that would work on the resulting structure.

Bibliography

[Allen 79] Allen, James.
A plan-based approach to speech act recognition.
PhD thesis, University of Toronoto, 1979.

[Austin 69] Austin, J.L.
How to do things with words.
Oxford University Press, 1969.

[Beaugrande 80]
Beaugrande, Robert de.
Advances in Discourse Processes. Volume IV: *Text, Discourse and Process: Toward a Multidisciplinary Science of Texts.*
Ablex Publishing, 1980.

[Brachman 78] Brachman, R. J.
A structural paradigm for representing knowledge.
Technical Report, Bolt Beranek and Newman, Inc. , 1978.

[Brachman 79] Brachman, R. J.
On the epistemological status of semantic networks.
In Findler, Nicholas (editor), *Associative Networks: the representation and use of knowledge in computers*, . Academic Press, 1979.

[Brachman 82] Brachman, Ronald J.
What "isa" is and isn't.
In *Proceedings canadian society for computation studies of intelligence.* Canadian Society for Computational Studies of Intelligence, 1982.

[Charniak 81] Charniak, Eugene.
The case-slot identity theory.
Cognitive Science 5(3):285-292, 1981.

[Chester 80a] Chester, Dan.
Using HCPRVR.
Technical Report, The University of Texas at Austin, 1980.

[Chester 80b] Chester, D.
HCPRVR: a logic program interpreter in LISP.
In *Proceedings of AAAI*. AAAI, 1980.

[Cohen and Perrault 79]
Cohen, Phillip R. and Perrault, C. Raymond.
Elements of a plan-based theory of speech acts.
Cognitive Science 3(4):177-202, 1979.

[Cohen, Perrault, and Allen 81]
Cohen, P. R., Perault, C. R., and Allen, J. F.
Beyond Question-Answering.
Technical Report, Bolt Beranek and Newman Inc., 1981.

[Correira 80] Correira, Alfred.
Computing story trees.
American Journal of Computational Linguistics 6:135-149, 1980.

[Cullingford 78]
Cullingford.
Script application: computer understanding of newspaper stories.
PhD thesis, Yale University, 1978.

[Dictionary 58]
Friend, Joseph H. and Guralnik, David B. (editors).
Webster's New World Dictionary.
The World Publishing Company, 1958.

[Fahlman 77] Fahlman, Scott E.
Netl; a system for representing and using real-world knowledge.
Mit Press, 1977.

[Fillmore 68] Fillmore, Charles.
The Case for Case.
In Bach, Emmon & Harms, Robert (editors), *Universals in Linguistic Theory*, . Holt, Rinehart, & Winston, 1968.

[Grimes 75] Grimes, Joseph E.
The thread of discourse.
Mouton Publishers, 1975.

[Grosz 77] Grosz, B.J.
The Representation and use of focus in dialog understanding.
PhD thesis, University of California at Berkeley, 1977.

[Hendrix 75] Hendrix, Gary G.
Expanding the utility of semantic networks through partitioning.
In *Advance Papers of the International Joint Conference on Artificial Intelligence.* SIGART, 1975.

[Hendrix 77a] Hendrix, G.
Human engineering for applied natural language processing.
In IJCAI-77 (editor), *Proceedings of the fifth international joint conference on artificial intelligence.* IJCAI, 1977.

[Hendrix 77b] Hendrix, G. G.
Lifer: a natural language interface facility.
Sigart Newsletter (61):25-26, 1977.

[Hendrix 79] Hendrix, Gary G.
Encoding knowledge in partitioned networks.
In Findler, Nicholas (editor), *Associative Networks: the representation and use of knowledge in computers*, . Academic Press, 1979.

[Hendrix, Sacerdoti, Saglowicz, and Slocum 78]
Hendrix, G. G., and Sacerdoti, E. D., Sagalowicz, D. and Slocum, J.
Developing a natural language interface to complex data.
ACM Transactions on Database Systems 3(2):105-147, 1978.

[Hobbs 77] Hobbs, Jerry.
Coherence and interpretation in english texts.
In *Proceedings of the fifth international joint conference on artificial intelligence.* IJCAI, 1977.

[Hobbs 79] Hobbs, Jerry.
Coherence and coreference.
Cognitive Science 3(1):67-90, 1979.

[Kintsch and van Dijk 75a]
Kintsch, W. and van Dijk, T. A.
Comment on se rappelle et on resume des histoires.
Langage :98-116, 1975.

[Kintsch and van Dijk 75b]
Kintsch, W. and van Dijk, T. A.
Recalling and summarizing stories.
1975.
Unpublished english translation of "Comment on se rappelle et on resume des histoires".

[Lehnert 77] Lehnert, Wendy.
The Process of Question Answering.
PhD thesis, Yale University, 1977.

[Lehnert 81a] Lehnert, Wendy.
Affect and memory representation.
In Third annual conference of the cognitive science society. Cognitive Science Society, 1981.

[Lehnert 81b] Lehnert, Wendy.
Plot Units and Narrative Summarization.
Cognitive Science 5(4):293-331, 1981.

[Levine 80] Levine, Sharon.
Questioning english text with clausal logic.
Master's thesis, The University of Texas at Austin, 1980.

[Miller 72] Miller, George A.
English verbs of motion: a case study in semantics and lexical memory.
In Melton, Arthur W. and Martin, Edwin (editors), Coding processes in human memory, . V. H. Winston & Sons, 1972.

[Miller & Johnson-Laird 76]
Miller, George & Johnson-Laird, Phillip.
Language and Perception.
Cambridge: Harvard University Press, 1976.

[Minsky 75] Minsky, Marvin.
A framework for representing knowledge.
In Winston, Patrick (editor), The psychology of computer vision, . McGraw-Hill, 1975.

[Parnas 72] Parnas, D. L.
On a criteria to be used in decomposing systems into modules.
Communications of the ACM 15(12):1053-1058, 1972.

[Protter 61] Protter, Eric (editor).
A children's treasury of folk and fairy tales.
Channel Press Inc, 1961.

[Quillian 68] Quillian, M. Ross.
Semantic Memory.
In Minsky, Marvin (editor), *Semantic Information Processing*, . Mit Press, 1968.

[Rich 79] Rich, Elaine.
User modeling via stereotypes.
Cognitive Science 3(4):329-354, 1979.

[Rumelhart 75] Rumelhart, David E.
Notes on a schema for stories.
In Bobrow, Danial G. and Collins, Allan (editors), *Representation and Understanding*, . Academic Press, Inc, 1975.

[Schank & Abelson 77]
Schank, R. C. and Abelson, R. P.
Scripts, plans, goals, and understanding.
Lawerence Erlbaum Associates, 1977.

[Schubert 79] Schubert, Lenhart K., Goebel, Randolph G., and Cerconce, Nicholas J.
The structure and organization of a semantic net for comprehension and inference.
In Findler, Nicholas (editor), *Associative networks: the representation and use of knowledge in computers*, . Academic Press, 1979.

[Scragg 75] Scragg, G. W.
Answering questions about processes.
In Norman, Donald and Rumelhart, David (editors), *Explorations in cognition*, . W. H. Freeman and Co., 1975.

[Searle 69] Searle, J. R.
Speech Acts.
Cambridge University Press, 1969.

[Searle 76] Searle, J. R.
A taxonomy of illocutionary acts.
In Gunderson, K. (editor), *Language, Mind and Knowledge*, . University of Minnesota Press, 1976.

[Simmons 82] Simmons, Robert F.
A narrative schema in procedural logic.
In Clark, K. L. and Tarnlund, S.-A. (editor), *Logic programming*, . Academic Press, 1982.

[Simmons 83] Simmons, Robert F.
Computations from the english.
Prentice-Hall, 1983.
In Press.

[Simmons & Correira 80]
Simmons, Robert F. and Correira, Alfred.
Rule forms for verse, sentences, and story trees.
In Findler, Nicholas (editor), *Associative networks: the representation and use of knowledge in computers*, . Academic Press, 1980.

[Smith 82] Smith M. K.
Knowledge-Based contextual reference resolution for text understanding.
PhD thesis, The University of Texas at Austin, 1982.

[van Dijk 76] van Dijk, T. A.
Macro-structures and Cognition.
In *Twelth annual carnigie symposium on cognition*. Carnigie Foundation, 1976.

[van Dijk 77] van Dijk, Teun A.
Text and context: exploration in the semantics and pragmatics of discoure.
Longman, 1977.

[Waltz 78] Waltz, David.
An english language question answering system for a large relational database.
Communication of the ACM 21(7):526-539, 1978.

[Wilensky 78] Wilensky, R.
Understanding goal-based stories.
PhD thesis, Yale University, 1978.

[Wilensky 81] Wilensky, R.
PAM.
In (editor), Schank, R. C. and Riesbeck, Christopher (editors), *Inside computer understanding: five programs plus miniatures*, . Lawrence Erlbaum Associates, 1981.

[Winograd 72] Winograd, Terry.
Understanding natural language.
PhD thesis, MIT, 1972.

[Woods 75] Woods, William A.
What's in a link: foundations for semantic networks.
In Bobrow, Daniel G. and Collins, Allan (editors), *Representation and Understanding*, . Academic Press, Inc., 1975.

[Woods, Kaplan, and Nash-Webber 72]
Woods, W. A., Kaplan, R., and Nash-Webber, B.
The lunar sciences natural language information system: final report.
Technical Report, Bolt Beranek and Newman, Inc., 1972.