

## How to Learn/What to Learn

Roger C. Schank and Mallory Selfridge  
Department of Computer Science  
Yale University  
New Haven, Connecticut 06520

**Abstract:** This paper discusses the kind of information that must be present in a computer program that models the linguistic development of a child. A three stage model is presented that characterizes the development of a natural language parser in a child of ages one, one and a half, and two. Some data from children of these ages is presented. General problems with respect to computer learning are also discussed.

**keywords:** learning, child development, parsing, conceptual analysis, psychological modelling.

### 1. Introduction

Learning has always been a magical word in AI. As long as there has been research in AI, there has been research on learning as well (Selfridge (1959), Samuel (1963), Winston (1970), and Sussman (1975) being some notable examples). There seem to have been two approaches (not necessarily mutually exclusive) with respect to the learning issue; those who felt that any program they were building would be better off if it had some ability to learn new rules or new material (e.g. Samuel (1963)) and those who felt that learning was a problem best pursued after we had a good idea of what it was that people knew in the first place.

For the last few years we have been developing conceptual representations for natural language information. These representations have been augmented recently to include larger knowledge structures such as scripts, plans, goals and themes (Schank and Abelson (1977)). We believe that we have a handle on what kinds of representational structures for knowledge need to be learned, and we assume an adult understander can be effectively characterized in terms of such structures. The question is how are such structures and representational systems learned?

To answer this question effectively, we believe that it is necessary to study how children learn. Children are, after all, processors that seem to start with almost nothing, and become highly skilled language users.

This work was supported in part by the Advanced Research Projects Agency of the Department of Defense and monitored under the Office of Naval Research under contract N00014-75-C-1111

The particular domain of learning we are interested in is language acquisition. There is reason to believe that a child begins the process of learning language very early on. There is evidence that a child of one year has all of the primitive ACTs of Conceptual Dependency available for representing his knowledge (Schank (1973)). Thus a child can begin learning a language equipped with predictions about the meaning of what he hears.

Let us consider a simple example of the learning process: Hana Schank, at age one was asked to "put your finger in your ear". She did so. As far as we can tell it was the first time she ever responded to the word 'put' correctly. (If it was not the first time, then there was some other first time. The important point being that there has to be one first time). How did she do it?

We advance the following hypothesis. She knew the words finger and 'ear'. She knew she was being asked to do an action (perhaps because a large number of the sentences that adults direct toward children of this age require them to respond by doing something.) She also knew and had used all the primitive ACTs and conceptual cases and relations in Conceptual Dependency (see Schank (1973a) for the argument about why this can be said to have been the case). She thus was able to reformulate the request into "pick a primitive ACT that has self as actor and can take finger and ear in one of its conceptual cases". This conforms to asking "What can I do with an ear and with a finger?" The answer to the first is 'nothing' (a child this age knows nothing about the physical act of hearing and since it has never been asked to put its ear to anything is unlikely to think of doing this without explicitly being shown how). The answer to what can you do with a finger is the primitive act MOVE. Thus the meaning of the request is a conceptualization which has 'Hana' as the actor, MOVE as the primitive act, 'finger' as the object, and in which the source and goal of the directive case are unknown. It is clear where 'ear' must fit in this conceptualization. It cannot be the source of the MOVE since 'finger' is not currently located at 'ear', so it must be the goal of the directive case.

Thus the action is correctly performed regardless of her knowing the actual word 'put'. Furthermore, such correct actions and their accompanying rewards cause the word 'put' to become a candidate for meaning MOVE. (Actually 'put' means far more than MOVE, but this is a good initial hypothesis). What this really means is that Hana already knew about 'putting' before learning the word for it, as discussed above and in (Schank (1973a)), and that this knowledge is exploited when learning the word.

The strategy employed by Hana has been effectively modelled in the SAM system by Granger (1977) in his FOUL-UP program. FOUL-UP is a program that has a deeper, more adult knowledge

base than Hana has, but it exploits the same principle: predict the intent of a word based on the possible meaning it could take in context.

The above word meaning learning technique was gathered from our own observations of several children. In general, our ideas of what are interesting phenomena in language learning, and our ideas of the knowledge needed for such learning, have been formed from numerous informal sessions with children, rather than by formal experimentation. This does not mean that such theories need be hazy or ill-formed. Since we are interested in process models of language acquisition, our theories must be quite specific if they are to be programmable.

It would be nice if we could exploit the research going on in psychology on children's acquisition of language. However, research in child language development has not produced results which are useful to us. It is our belief that the fundamental issues in child language research are ones of the development of language comprehension (as opposed to generation), and the acquisition of world knowledge. Psychology, on the other hand, tends toward a somewhat different point of view. The emphasis is usually on the production of language, rather than its comprehension, and attention is frequently focussed on the development of the syntax of language, rather than the evolving and growing meanings the child is able to deal with. Furthermore, psychologists do not emphasize process models in their research, as AI does do, but are interested primarily in descriptive models. Thus much current psychological research in language acquisition is not directly relevant to our interests and goals.

Dale (1976), for example, in a review of recent work in language development, discusses at considerable length several theories of syntax acquisition in generation, but considers only relatively briefly the development of meaning in generation and understanding. He does not consider the role of an internal world model in the child's language learning. Brown (1973) and Halliday (1975) are examples of psychologists who indeed do present theories with semantic emphasis, but both are handicapped by the lack of an adequate meaning representation and by an inadequate notion of process. Macnamara (1972) comes closest to the philosophy of the authors, in that he suggests that an infant learning language already knows, or can infer, much of the meaning of things that are said to him, and that his problem is to figure out how that meaning is encoded in the utterances he hears. Macnamara doesn't discuss the kinds of 'available meanings' that a child might have, however, nor how a language-meaning relationship is constructed.

## II. What to Learn

In order to effectively approach the problem of learning in children it is a good idea to isolate stages of development and write mapping rules to get us from one stage to another. In

this paper we will describe the stages that children can be naturally said to be at. These stages represent the development of increasingly sophisticated world knowledge on the part of the child. We will then discuss a set of programs which models parts of several of these stages, and which models aspects of the progression from one stage to the next.

When we talk of stages of development it is important to differentiate several kinds of stages. We will argue here that there are at least seven types of things that are learned by a small child. (We are by no means saying that a child learns only these seven types of things. Indeed, he learns considerably more than these. The things we will talk about are those items crucial to the understanding process in adults that we have modelled in our computer understanding programs.) What follows is a summary of each of these seven items, along with a brief estimation of a child's abilities at these items at several different ages. We will present no evidence backing up our assertions here since this is beyond the scope of this paper. Instead we refer the reader to Schank (1973), Schank and Abelson (1977), Nelson (1975), and Nelson and Gruendel (1977) for the details and data backing up some of these assertions.

**Conceptual Dependency Representation (CD):** This is the representation system for events. The basic construct is actor-action-object-direction groupings. There are eleven primitive actions and a small set of relations between concepts (Schank, 1975). By age one, a child can use all of the primitive acts, and by age three and a half can converse about mental acts like MBUILD.

**Causal Links:** These are the relationships between events. There are two physical links and two mental links. There are also a few abstract relationships (Schank, (1973b)). By age one, a child knows about physical causality, and by age three and a half can converse about mental causality.

**Conceptual Analyzer:** This conforms to our parsing program for adults (Riesbeck and Schank, (1976)). It is a procedure that maps English strings into CD. A child of age one uses keywords for parsing, and can, for example, point to objects. By age two, a child has added some CD, and can respond to commands. By age two and a half, the child is starting to use conceptual parsing, and can respond in a manner similar to adults. By age three and a half, he has understood when he knows something, and by age four is beginning to learn the rules for conversation.

**Scripts:** These are large, standard combinations of events that describe a well known situation such as a restaurant, going to bed, riding in a car (Schank and Abelson, (1977)). At age one, a child has formed simple scripts, but not until age two are they present for day to day use. By age two and a half scripts have been formed for all known situations, though a lack of informa-

tion about some situations leaves a child's script base incomplete at least through age three and a half. By age four, a child has basically an adult's scriptal abilities, though still lacking many particular scripts.

**Plans:** These are the general, goal-driven mechanisms that generate scripts (Schank and Abelson, (1977), Meehan (1976)). These develop slowly in a child. They are just starting to have significance by age three and a half, and are being used in understanding a little by age four.

**Generator:** This is the mechanism for mapping concepts into sentences. A child can generate what are words to adults by age one or so. By age two four word sentences with action and object can be generated, and by two and a half or so paragraphs with short sentences are possible. When a child is three and a half and four his generation is nearly adult.

**Memory:** This is the storage mechanism for facts. The form of memory is much debated, however we can safely assume children to have one. See Rieger (1975), Quillian (1968), Norman and Rumelhart (1975) for various conceptions of what memory is like. In our terms, a child of age two, and possibly of age one, has episodic memory. By age two and a half or so, scriptal memory has been added, and by age three and a half a rudimentary semantic memory exists.

We view a child as developing along all (and more) of these dimensions simultaneously. Of course, the completed development of one facet of the child is often a prerequisite for the beginning of another. This is why Conceptual Analysis waits for CD and Causal Links to be completed. Similarly, generation cannot begin seriously until the script-based memory processes have given the child the ability to organize its experiences and thus have something to say about its experiences.

It is our ultimate aim to simulate a child on the computer. That is, we would like to have a program which possesses certain innate abilities (the ability to grasp, the desire to move new objects to its mouth, to cry when hungry, for example), and which lives in a simple environment such as that a child lives in. We would like to explore how learning can be modelled through interactions between the innate abilities and the environment to produce a program which develops the ability to understand. This is, of course, a tremendously ambitious project. Our first pass at it, then, will be to develop only one aspect of the child while faking the other developmental parts that are interdependent on that part. In this paper therefore, we will discuss only the development of the conceptual analyzer.

### III. Learning to Understand

We will divide up the process of the acquisition of a conceptual analysis capability (hence parsing) into three distinct stages. Naturally the development of parsing, like anything else, is really a continuous process, rather than one that proceeds by discrete stages. Nonetheless, it is useful to maintain the stage division idea. This section describes a three stage model of the language acquisition in a child during his second year. The data for this model and the programs implementing it were obtained from observations of several children during the ages one to two.

#### Stage 1

At the end of the first year of a child's life, he can usually understand (and respond to) questions of the form: "Where is X?"; where X is a part of his body, a toy, a room, or the name of a person. His responses in the above four cases are: point to bodypart; get or point to the toy; go or point to the room; go or point to the person. The parser that is necessary to perform this task is relatively trivial. It requires knowing the name of the object and knowing the semantic relationships that determine the appropriate actions. The parsing part requires no more than recognizing an intonation pattern denoting a question, recognizing the name of the object, and determining from past experience what actions are associated with that object that the child could possibly be being asked to perform.

CHILD1 is intended to be a model of a one year old's parsing ability. It accepts input sentences similar to those which might be addressed to a one year old, and prints out a conceptual dependency structure representing the behavior it is generating in response. It has a simple set of test-action rules, which build the structure, and a simple monitor which administers the rules. The monitor looks at the words in the input sentence one at a time, and then sees whether there is a rule with a test that applies to that word in that situation. If it finds one, it executes its action. The successive execution of the actions of the rules builds the CD structure.

The test of a rule requires taking several things into account. In CHILD1, the rules test to see whether the word refers to a bodypart, a toy, a room, or a person, and whether the toy or person is near or far from where the program 'is' and whether, if far, whether the program is in the same room as the toy or person.

The rules which model this stage are written in LISP, as are the rules for the other two stages, and below are provided English statements of the rules. Note that CHILD1 has a short-term memory capacity of only one known word.

- R1) if a word refers to an object or place and no such word has been seen yet then load it into the short-term memory where it can be examined by other rules.
- R2) if a CD structure has not been built and the word refers to an object which is 'close' then construct the CD structure corresponding to pointing to the object, pointing to the object.
- R3) if a CD structure has not been built and the word refers to an object with is 'far' then construct the CD structure corresponding to going to where the object is.

Below is a sample interaction with CHILDi. The user's input is in lower-case letters, and follows a left-arrow prompter. CHILDi responds by printing the conceptual dependency diagram corresponding to its action following the input.

```
>go get the ball
((ACTOR (CHILD) ACT (MOVE) OBJECT (FINGER)
TO (BALL) FROM (NIL)))

>where is Roger ?
((ACTOR (CHILD) ACT (PTRANS) OBJECT (CHILD)
TO (ROGER) FROM (NIL)))

>where is Mallory ?
((ACTOR (CHILD) ACT (MOVE) OBJECT (FINGER)
TO (MALLORY) FROM (NIL)))

>put the ball on the table
((ACTOR (CHILD) ACT (MOVE) OBJECT (FINGER)
TO (BALL) FROM (NIL)))

>get the ball from in the kitchen
((ACTOR (CHILD) ACT (MOVE) OBJECT (FINGER)
TO (BALL) FROM (NIL)))
```

Note that no attention is paid to 'where', or 'put'. CHILDi is only prepared to respond to words which signify things in the environment which it has behaviors to operate on. At this stage the same is true of children, who are happy to respond to "Frimble Mama" by pointing to Mama if the intonation is question-like.

## Stage 2

By age one and a half a child can process sentences of the order of complexity discussed in the beginning of this paper. The child is learning words such as 'put' by relying on the primitive ACTs and the relationships that hold between them and their possible conceptual uses. Thus, names of objects are learned first, but names of actions can be learned by associating unknown words with the action that the child has determined is the one he is being asked to perform. The child is learning to rely more heavily on the conceptual case requirements of the actions he is performing. Thus, at the same time, he is also learning to handle two objects

at a time in one sentence. That is, he can also determine what possible relationship could hold between two objects and perform it. Sentences such as "put the ball in the box" can be correctly handled by the child precisely because the relationship he knows to hold between balls and boxes happens to be the one that is being requested of him. At this stage, "put the ball near the box" is misunderstood as being identical with the sentence above. (This is due to the fact that 'in' relationships are usually more important than 'near' relationships in a child's world and are more likely to have been emphasized by his parents.)

The model that accounts for this stage is a parser that is substantially the same as in the first stage, but with the exception that more than one object can now be handled. Thus, the child more fully exploits the case requirements of the primitive action that he has selected to act out. PTRANS, for example, when selected as being what you ordinarily do with a ball, requires a directional case which can be used to predict where other objects in the sentence might fit in relation to the described action. Thus, in the sentence "ball Papa", if PTRANS is the act, then the fact that the object of the PTRANS must be 'PTRANS-able' implies that the filler of the OBJECT slot is 'ball', not 'Papa'. Since the ACTOR must (at this stage) be the child, only the TO and FROM slots remain to be filled. Since the ball is (presumably) near the child, it must be the TO slot that is filled with 'Papa'. Unknown words, repeatedly used in conjunction with these inferred primitive ACTs, are learned as tags to the primitive ACTs directly. Thus, in the sentence "put the cup on the table", the act PTRANS can be inferred from what can be done with a cup and a table, and this information can then be attached to the unknown word 'put'.

The program which models the child at age one and a half, CHILDi.5, is structurally the same as CHILDi. It uses the same monitor, and also has a set of rules which produce actions. The set of rules used by CHILDi.5 includes the same rules as those used by CHILDi, plus several more. These allow for the construction of executable structures by taking into account constraints on what can fill what slot. These constraints are imposed by semantic information attached to words. In addition, CHILDi.5 has a rule which allows it to take a conceptual structure which has been built by examining the known words in a sentence, and hypothesize that this structure is part of the meaning of other words in the sentence which it does not know. Other rules, not yet a part of CHILDi.5 but under development, will be used to refine such hypothesized meanings in order to specify precisely what conceptual structure should be built when a particular word is heard. CHILDi.5 can deal conceptually with up to two words which refer to things or which have default actions associated with them. That is, it can accommodate only two objects or words with default actions into its structure building process, and further such

words are ignored. Its rules (in addition the rules of CHILD1) are as follows.

- R4) if the input word is unknown, then save it for use by R9.
- R5) if the word has a primitive act associated with it, then build the structure specified by the act and activate demons for filling slots in that structure.
- R6) If one input word is associated with a default primitive act, and another input word refers to an object which fills a slot of that act's structure, then build that structure and fill that slot with that object.
- R7) if the first word seen can build a structure which has a slot for the object referred to by the second word, then build that structure.
- R8) if the second word seen can build a structure which has a slot for the object referred to by the first word, then build that structure.

Rules R6, R7, and R8 implement the following idea: if there were two thing-words in the input, and one of them can have something done to it which involves the other, then that action must be the meaning of the input, and that structure is built. The rule below is a brief verbal description of CHILD1.5's learning rule.

- R9) if an unknown word has been heard, and a CD structure has been actively (not by default) constructed, then 'hypothesize' that the filler of the ACT slot is associated with the unknown word by storing that act under the new word along with slot-filling demons constructed by examining the fillers of the existing structure.

The net result of this rule is to take an unknown word and associate with it information which will allow it to be used by R5 to construct a CD structure for the input. Below is a sample session with CHILD1.5. Note that CHILD1.5 can integrate the meanings of several different words, but does not yet fully understand the meaning of sentences. In particular, it does not understand prepositions, but merely uses context and 'common sense'. Note also that this session demonstrates CHILD1.5 learning a meaning for the word 'put'. It knows that the words 'hat' and 'rack' refer to objects, but it does not have any default actions associated with them. So, it uses one of CHILD1's rules to generate a response. After learning what act 'put' refers to, by means of R9, above, it then can construct a meaning for 'put the hat on the rack'.

```
>the cup is on the table
((ACTOR (CHILD) ACT (PTRANS) OBJECT (CUP)
TO (TABLE) FROM (NIL)))
```

```
>put the hat on the fribble
((ACTOR (CHILD) ACT (PTRANS) OBJECT (CHILD)
TO (HAT) FROM (NIL)))
```

```
>put the cup on the table
((ACTOR (CHILD) ACT (PTRANS) OBJECT (CUP)
TO (TABLE) FROM (NIL)))
```

```
>put the hat on the rack
((ACTOR (CHILD) ACT (PTRANS) OBJECT (HAT)
TO (RACK) FROM (NIL)))
```

### Stage 3

By age 2 the simple program of Stage 2 is expanded into a more complex one. Children of age 2 can respond to the following kinds of sentences:

Put the X in the Y	Get your X
Do you want to Z?	Give me an X
Are you in 'some state'	Go into the Y

Clearly the parser of the year and a half old child has completed its work and many verbs have been learned. It is tempting to argue at this point that a two year old child has a normal adult-like conceptual analyzer. However, we would prefer a less drastic progression from the one and a half year old parser to the two year old one.

The argument we will make is that the two year old is a "key concept" parser rather than a key word parser. What distinguishes him from the one and a half year old is his knowledge of verbs and their associated prepositions, and his ability to use the primitive act associated with a verb in order to predict slot instantiation in the conceptual structure being built. (Note that prepositions can be learned in a manner analogous to that described for verbs above. That is, when there are relationships between objects which the child has inferred must exist, unknown words in the input sentence are good candidates for names for these relationships. Certainly other factors enter in, however. In order for a child to learn words other than 'in' in the ball and box example, it is necessary to repeat again and again to a child "no, NEAR the box". This must also be accompanied by showing the child what exactly to do.) Thus the two year old parser is a top down predictive analyzer (as is ELI, our current parser). When a verb is recognized the cases are predicted for it and objects present in the sentence are put into their correct conceptual place in accordance with the prepositions found. If the verb is not recognized, or the prepositions are not recognized, then the two year old reverts to the old techniques that worked before when he was younger.

Before we present the program implementing our two year old model, we will present some data from Joshua Schank, age 2 years (exactly). First a context is presented, and then Joshua's reac-

tion in that context. Joshua's reaction is

we are doing puzzles in the study,  
"there is one more piece, get that  
piece"

look around

"Go get a book from Hana's room"  
go to Hana's room, get a  
book, bring it back.

Josh is playing with the vacuum cleaner,  
"are you getting the floor clean?"  
ignore

"Get some brownie for Mallory, Josh.  
Get some brownie for Mallory"  
Walk to where brownies are  
on counter, stop, note that  
they can't be reached, look  
at father.

Roger gives brownie to Josh, "()K,  
go give the brownie to Mallory."  
Run to Mallory and give  
him the brownie.

"Want to color? Get me a piece of paper."  
Brings a piece of paper  
to Mallory

"Hey Josh, get the chair, get the chair."  
Turn around, get a chair,  
bring it closer to the table,  
table. Say "Hey, that's  
better", get on it.

"Josh? get ON the chair, get ON the  
chair, get on the chair."  
Climbs into the chair.

"Get me the tape recorder, Josh"  
Go get tape recorder and  
brings it to Mallory

"Get on the tape recorder, Josh"  
"get on the tape recorder."  
"get on the tape recorder."  
Looks at Mallory with a  
funny expression. Go over  
to tape recorder, put foot  
next to it (as if to step  
on it).

"go get your cup"  
Go to Mama and say "I want  
my cup."

MS holds up a small toy barrel and says  
"where did you get this?"  
say "barrel"

"Get me a box of tissues, Josh"  
run into playroom from  
kitchen say "box of  
tissues in here?"

The program, CHILD2, which models the two year old child again uses the monitor used in the previous two programs, and only differs in that new rules have been added and definitions of verbs have been included. The rules which have been added are ones that deal with prepositions, constructing a conceptual structure encoding the prepositional phrase (the conceptual structure actually constructed is currently only shorthand.) The definitions are sets of expectations and default specifications. Both are implemented as procedures which either look for a certain kind of word to fill a slot on the action structure, or which fill an unfilled slot with a default value if that slot was not filled by something in the input. Below are the preposition-handling rules and an example verb definition.

RIO) if the word is a preposition  
then save it, and load RII.

RII) when a preposition is being saved, and the current word refers to a thing or place, then construct a conceptual representation for the place denoted by the preposition and the object, and put that representation into short-term memory, accessible to other rules.

Rules RIO and RII construct simple meaning representations for prepositional phrases. Below is shown an English rendition of CHIL3's definition of the word 'give'.

definition of: GIVE

structure built: ACTOR (CHIL3)  
ACT (PTRANS)  
OBJECT (objl)  
TO (recipl)  
FROM (place 1)

filler constraints: objl must be a small object  
recipl must be a person  
placel must specify a place

This definition is used when the word 'give' is found in the input. SeeinR 'give' causes the specified structure to be built, and causes requests to be activated which look for words referring to objects which meet the filler constraints.

Below is presented a sample session with CHIL2. Note that it knows several verbs, and that it has some knowledge of prepositions (although, again, the representation employed here for prepositions is hardly adequate).

>give me the cup  
((ACTOR (CHILD) ACT (PTRANS) OBJECT (CUP)  
TO (MALLORY) FROM (NIL)))

>get the cup on the table  
((ACTOR (CHILD) ACT (PTRANS) OBJECT (CUP)  
TO (NIL) FROM (ON-TABLE)))

>get me the ball from the kitchen  
 ((ACTOR (CHILD) ACT (PTRANS) OBJECT (BALL)  
 TO (MALLORY) FROM (KITCHEN)))

>get on the chair  
 ((ACTOR (CHILD) ACT (PTRANS) OBJECT (CHILD)  
 TO (ON-CHAIR) FROM (NIL)))

>did Roger give the cookie to Diane ?  
 ((ACTOR (CHILD) ACT (PTRANS) OBJECT (COOKIE)  
 TO (DIANE) ))

>get me on the chair  
 ((ACTOR (CHILD) ACT (PTRANS) OBJECT (SOME-OBJECT)  
 TO (MALLORY) FROM (ON-CHAIR)))

Note that CHILD2 by no means understands everything of what is said to it. It does, however, make mistakes which are the sort that children make, and it constructs plausible meanings given the information it has.

#### IV. Conclusion

We have attempted to show in this paper that it is possible to model developmental stages of a child on a computer. Rules that map one stage into another which make minimal assumptions (such as "somewhere between age one and one and a half a child begins to be able to construct conceptualizations in which more than one conceptual relationship obtains between the referents of the words in the sentence") could be used to get the computer model to develop automatically, as a child does. CHILD1.5 contains a preliminary sample of such a rule.

We have two aims in this research. First, we are interested in how children learn to understand and we believe that building a computer model of such a process can shed light on questions regarding what knowledge must be innate (very little we would suspect) and what knowledge is learned.

But we also have another aim. We, like most AI researchers, wish to build a learning machine. We believe that the most promising approach is to model natural developmental stages (as exhibited by children), and that a learning machine will use processes similar to those employed by children.

A learning model must start in the right place and make developmental sense. The CHILD programs are only a very small start in that direction. We hope to uncover, by stepping our CHILD through various stages of development, the learning principles that people actually employ.

#### References

- 1] Brown, R. (1973), *A First Language / The Early Stages*. Harvard University Press, Cambridge, Ma.
- 2] Dale, P. S. (1976), *Language Development*. Holt Rinehart and Winston, New York.
- 3] Granger, R. H. (1977). *FOUL-UP: A Program That Figures Out Meanings of Words from of Words from Context*. Fifth International Joint Conference on Artificial Intelligence, August 1977, Cambridge, Ma.
- 4] Halliday, M. A. K. (1975), *Learning How to Mean - Explorations in the Development of Language*. Edward Arnold, London.
- 5] Macnamara, J. (1972). *Cognitive Basis of Language Learning in Infants*. *Psychological Review*, 1972, 79, 1-13.
- 6] Meehan, J. (1976). *The Metanovel: Writing Stories by Computer*. Ph. D. thesis, Yale University.
- 7] Nelson, K. (1974). *Concept, Word, and Sentence: Interrelations in Acquisition and Development*. *Psychological Review*, 1974, 81, 267-285.
- 8] Nelson, K. and Gruendel, J. (1977). *Young Children's Memory for Real Life Events*. *Life Events*, paper presented at the Symposium on Cognition, Carnegie-Mellon University, May 1977
- 9] Norman, D. and Ruraelhart, D. (1975). *Explorations in Cognition*. W. H. Freeman and Co., San Francisco, Ca.
- 10] Quillian, M. R. (1968). *Semantic Memory*. In M. Minsky, ed. *Semantic Information Processing*. MIT Press, Cambridge.
- 11] Rieger, C. (1975). *Conceptual Memory*. In R. C. Schank, ed. *Conceptual Information Processing*. North Holland, Amsterdam.
- 12] Riesbeck, C. K. and Schank, R. C. (1976), *Comprehension by Computer: Expectation-Based Analysis of Sentences in Context*. Technical Report #78, Computer Science Department, Yale University, New Haven, Conn.
- 13] Samuel, A. L., (1963). *Some Studies in Machine Learning Using the Game of Checkers*. In Feigenbaum, E. A. and Feldman, J. (eds.), *Computers and Thought*, McGraw-Hill, New York.
- 14] Schank, R. C. (1973b). *Causality and Reasoning*. Technical Report #1, Istituto per gli Studi Semantici e Cognitivi, Castagnola, Switzerland.
- 15] Schank, R. C. (1973a). *The Development of Conceptual Structures in Children*. Stanford Artificial Intelligence Lab. Memo AIM-203.
- 16] Schank, R. C. (1975). *Conceptual Information Processing*. North Holland, Amsterdam.
- 17] Schank, R. C. and Abelson, R. P. (1977). *Scripts, Plans, Goals, and Understanding*. Lawrence Erlbaum Press, Hillsdale, N.J.
- 18] Selfridge, O. G. (1959) *Pandemonium: A Paradigm for Learning*. *Proceedings of the Symposium on Mechanisation of Thought Processes*. Eds. Blake and Uttley. H.M. Stationary Office, London.
- 19] Sussman, G. (1973), *A Computational Model of Skill Acquisition*. MIT-AI Technical Report 297, Cambridge, Ma.
- 20] Winston, P. (1970), *Learning Structural Descriptions From Examples*. MIT-AI Technical Report 231, Cambridge, Ma.