A Computer Model of Child Language Learning

Mallory Selfridge

Yale University

Abstract

A computer program modelling a child between the ages of 1 and 2 years is described. This program is based on observations of the knowledge this child had at age 1, the comprehension abilities he had at age 2, and the language experiences he had between these ages. The computer program described begins at the age 1 level, is given similar language experiences, and uses inference and learning rules to acquire comprehension at the age 2 level.

Introduction This paper describes a computer model of the development of comprehension abilities in a child, Joshua, between the ages of one and two years. The program begins with the kind of knowledge that Joshua had at age 1, when he understood no language, and learns to understand commands involving action, object, and spatial relation words at Joshua's age 2 level. It does so by being given the kind of language experiences Joshua had between the ages of 1 and 2, and making use of rules to 1) infer the meaning of utterances, 2) attend to words, and 3) learn language meaning and structure. The program passes through a reasonable developmental sequence and makes the same kind of errors that children make at intermediate stages.

This work suggests that language learning to the 2 year old level can be accounted for primarily by the learning of word meaning and structure, that world knowledge is crucial to enable the child to infer the meaning of utterances, and that children hear language in situations which enable them to perform such inferences. The success of the program in modelling Joshua's language development -- both its progression and its errors -- suggests that it embodies a plausible theory of how Joshua learned to understand language. While there are several aspects of the model which are unrealistic (for example, segmented input, no ambiguous words, no simultanious conceptual development), there is reason to believe that future work can sucessfully address these issues. Further details can be found in Selfridge (1980).

This paper first considers Joshua's initial state of knowledge at age 1, and then his comprehension abilities at age 2. It describes the kind of language experiences he had, and several kinds of learning rules which can account for Joshua's development. The computer program incorporating these observations and rules is described, and finally some conclusions are presented.

Joshua's Initial Knowledge The first component of a computer model of the development of Joshua's comprehension is Joshua's knowledge prior to his language learning. Observations like the following suggest that Joshua had considerable knowledge of objects, actions, spatial relations, and gestures at age 1 (ages are given in YEARS:MONTHS:DAYS):

> 0:11:19 Joshua and I are in the playroom. I build a few block towers for him to knock down, but he doesn't do so; rather, he dismantles them removing the blocks from the top, one at a time
>
> 1:0:16 Joshua and I are in the playroom. Joshua takes a toy cup, and pretends to drink out of it.
>
> 1:2 Joshua is sitting in the living room playing with a ball. I hold my hand out to him, and he gives me the ball.

The above observations show that Joshua knew the properties and functions of objects like blocks, cups and balls. He knew actions that could be performed with them, and various spatial relations that they could enter into. Finally, he knew that behavior can be signaled through gestures by other people. Thus, a language learning program must be equipped with this kind of knowledge.

Joshua's Comprehension Abilities at Age 2 At age 2, Joshua could respond correctly to commands with unlikely meaning and structure. His correct responses suggests full understanding of them. For example, consider the following:

> 2:0:5 We walk into the living room and Joshua shows us his slippers. His mother says "Put your slippers on the piano." Joshua picks up the slippers and puts them on the piano keys, looking at his mother. She laughs and says "Thats silly." Joshua removes the slippers.

The meaning of this utterance is unlikely since slippers do not generally go on piano keys, and piano keys don't generally have things put on them. His response suggests that he was guided by full understanding of the meanings of the words in "Put your slippers on the piano."

At age 2 Joshua also understood language structure, as the following example shows:

> 2:0:0 Joshua and I are in the playroom, my tape recorder is on the floor in front of me.

I say "Get on the tape recorder, Joshua". Joshua looks at me oddly, and looks at the the tape recorder. I repeat "Get on the tape recorder." Joshua moves next to the tape tape recorder. I once more repeat "Get on the the tape recorder." Joshua watches me intently, and lifts his foot up and slowly moves it over the tape recorder to step on it. I laugh and pull the tape recorder away.

It seems that Joshua understood "Get on the tape recorder" the first time I said it, and that his reluctance to comply reflected his knowledge that what I was asking was very unlikely. That is, Joshua understood that the tape recorder was the object to be underneath him, although this is unlikely given his experience with it. This, in turn, suggests that Joshua understood the structure of the word "on", namely, that the word whose meaning is the supporting surface follows "on". Thus a program modelling Joshua at age 2 must understand utterances using language structure.

Joshua's Language Experiences In the year between the ages of 1 and 2, Joshua experienced situations which allowed him to make inferences concerning the utterances he heard. In this section, three examples of such situations are given, and inference rules accounting for Joshua's response and attention to words are presented.

In the first example, I am using an utterance and simultaniously signalling the meaning of that utterance through gestures:

1:2:17 We are sitting in the living room, Joshua is holding a book. I look at Joshua, maintain eye contact for a moment, hold my hand out to him and say "Give me the book, Joshua." Joshua holds the book out to me.

In this situation, Joshua probably inferred that the meaning of "Give me the book, Joshua." was the same as that signalled by the gestures. The following rule captures this idea:

Gestural Meaning Inference
If an utterance is accompanied by gestures with associated meanings then infer that the the utterance means the same as the gestures.

Knowledge of object function and properties helped Joshua infer responses in other situations. In the following, Joshua used his knowledge that books can be opened in his response:

1:0:9 Joshua has a book in his hand, and is looking at it, turning it over, and examining it. His mother says "open the book, open the book..." Joshua opens the book. She says, "Good Joshua, good."

A rule summarizing this inference is the following:

Function/Property Inference
If an utterance is heard while interacting with an object then the meaning of the utterance involves a function or property of that object.

Parent speech to children posesses many attention-focussing characteristics (e.g. Newport, 1973). The following example is typical:

1:18:0 Joshua's father is trying to demonstrate that Joshua knows the names of the upstairs rooms, and has put a toy lawnmower in the bathroom. He says "Where is the lawnmower, Josh? Its in the BATHROOM. The LAWNMOWER is in the BATHROOM. BATHROOM!"

Joshua's attention to "bathroom" in this example can be explained by the following rule:

Attention Inference
If a word is emphasised, repeated, or said in isolaytion, then attend to it.

These are the kind of rules which I postulate enabled Joshua to infer the meaning of utterances from context, and attend to part of the utterance. The program must be equipped with such rules and must be given input in similar contexts.

Learning Rules This section will consider Joshua's learning of action, object, and relation words, and language structure. It presents accounts of how Joshua might have learned each of these. Most of the rules have their roots in the learning strategies proposed by Bruner, Goodnow, and Austin ('1956). One way Joshua learned the names of objects is by having them named for him, as in the following example:

1:0:0 Joshua is crying. His mother picks him up and goes over to the refrigerator. She gets some juice, holds it up, and asks, "Do you want some JUICE?" Joshua keeps crying. She gets a banana and asks, "Do you want some BANANA, Joshua?" Joshua reaches for it.

The following rule models Joshua's ability to learn by having objects named:

Direct Naming Inference
If a word and an object are both brought to attention, infer the word is the object's name.

This rule, and other object word learning rules, can account for how Joshua learned object words such a "slippers", "piano", "ball", and "table".

Action words can be learned via inferences about other known words in the utterance. In the following example, summarized from Schank and Selfridge (1977), Hana could have inferred the meaning of "put" based on her knowledge of the meanings of "finger" and "ear."

(age 1) Hana knows the words "finger" and "ear", but not "put." She was asked to "Put your finger in your ear," and she did so.

The following two rules can account for learning "put" in situations like this. The first suggests that "put" would initially be learned as "put something in something else." The second, applied after the first in a slightly different situation, would refine the meaning of "put" to "put something someplace".

Response Completion Inference
Infer the meaning of an unknown word to be the meaning of the entire utterance with the meanings of the known words factored out.

Meaning Refinement Inference
If part of the meaning of a word is not part of the meaning of an utterance it occurs in, remove that part from the word's meaning.

Rules like the above can account for Joshua learning action words like "put", "bring", "give", and so on. However, they can also account for Joshua learning relation words, such as "on" and "in". If Joshua knew "put", "ball", and "box", say, and was asked to "put the ball in the box", these rules would account for his learning that "in" referred to the "contained" relation.

These, then, are the sort of rules the program uses to learn word meanings. The program's rule for learning language structure is more direct. It is based around the two structural predicates, PRECEDES and FOLLOWS, which relate the positions of words and concepts in short-term memory. This rule models Joshua's acquisition of structural information upon hearing utterances he understands, and appears below:

Structure Learning Rule
If a slot filler occurs preceding or following a word's meaning then update the word's definition that information.

This rule accounts for Joshua learning that the filler of the VAL slot of "in"'s meaning -- (CONT VAL (NIL)) -- is found FOLLOWing "in" in the utterance.

The Program This section presents four excerpts from a run of the program, written in LISP on a DECSYSTEM-20. Each represents the program at a different stage in development as it progresses from Joshua's age 1 abilities to Joshua's age 2 abilities, using the inference rules described previously. The knowledge representation used is Conceptual Dependancy (Schank, 1973), and the language understanding process embedded in the program is similar to that in Birnbaum and Selfridge (1979).

The first stage of the program corresponds to Joshua at age 1. At this stage, the program had only the knowledge ascribed to Joshua at that age. In the excerpt below, the "parent" types a lowercase utterance to the program, and the program responds with a message stating its lack of understanding. When the parent provides gestures via simulated visual input, however, the program understands, and prints the CD representation of its response.

```
|PARENT SAYS: give me the ball
|
|CHILD STARES BLANKLY AT PARENT
|CHILD RETURNS TO PLAY
|
|
|CHILD SEES: (PARENT HOLDS OUT HAND)
|        (PARENT LOOKS AT BALL1)
|
|CHILD INFERS RESPONSE USING RULE:
|        GESTURAL MEANING
|CHILD RESPONDS:  (ATRANS ACTOR (CHILD)
|   OBJECT (BALL1) TO (POSS VAL (PARENT))
```

In the second stage, shown in the excerpt below, the program has learned the meaning of several words, and understands some utterances correctly. In this case, it has learned the words "put", "ball", and "box". However, notice that although it responds correctly to the first utterance given by the parent, it misunderstands the second. This sort of error is reported in Hoogenraad et al. (1976). Not knowing "on", the program incorrectly infers that the appropropriate relationship is containment.

```
|PARENT SAYS: put the ball in the box
|
|CHILD INFERS RESPONSE USING RULE:
|     UTTERANCE UNDERSTANDING, FUNCTION/PROPERTY
|CHILD RESPONDS: (PTRANS ACTOR (CHILD)
|          OBJECT (BALL1) TO (CONT VAL (BOX1))
|
|PARENT SAYS: put the ball on the box
|
|CHILD INFERS RESPONSE USING RULE:
|    UTTERANCE UNDERSTANDING, FUNCTION/PROPERTY
|CHILD RESPONDS: (PTRANS ACTOR (CHILD)
             OBJECT (BALL1) TO (CONT VAL (BOX1))
```

The transition from the second stage to the third is accomplished by teaching the program more words. In this case it has learned the additional words "slippers", "on", "piano", "ball", and "table." At this stage, the program can now understand "Put the slippers on the piano", whereas at any earlier stage it would not have. The program also prints out a message showing that it recognizes this as an unusual request.

However, although this stage represents Joshua's age 2 understanding of word meaning, the program has not yet learned language structure. The program interprets the second utterance incorrectly, however, in accord with its knowledge of the usual relationships between objects. This sort of error is similar to that reported in Stroher and Nelson (1974).

```
|PARENT SAYS: put the slippers on the piano
|
|CHILD LOOKS AT PARENT STRANGELY
|
|CHILD INFERS RESPONSE USING RULE:
|      UTTERANCE UNDERSTANDING
|CHILD RESPONDS:
|  (PTRANS ACTOR (CHILD)
|  OBJECT (SLIPPERS1) TO (TOP VAL (PIANO1)))
|
|PARENT SAYS: put the table on the ball
|
|CHILD INFERS RESPONSE USING RULE:
|   UTTERANCE UNDERSTANDING
|CHILD RESPONDS: (PTRANS ACTOR (CHILD)
|   OBJECT (BALL1) TO (TOP VAL (TABLE1)))
```

The fourth stage is shown in the excerpt below. The program has now learned the structure of "on", and can hence correctly understand "Put the table on the ball." In addition, it prints out

a message indicating its awareness of the peculiarity of this request.

```
|PARENT SAYS: put the table on the ball
|
|CHILD LAUGHS AT UNUSUAL REQUEST
|
|CHILD INFERS RESPONSE USING RULE:
|   UTTERANCE UNDERSTANDING
|CHILD RESPONDS: (PTRANS ACTOR (CHILD)
|    OBJECT (TABLE1) TO (TOP VAL (BALL1)))
```

At the fourth stage, the program has successfully learned to understand a subset of language at Joshua's age 2 level. It began with world knowledge similar to that Joshua began with, was equipped with reasonable learning and inference rules, and progressed as he did by being given language experiences similar to those he experienced.

Conclusions This paper has described a computer model of a child learning to understand commands involving action, object, and relation words. The program learns language meaning and structure to the level attained by the child at age 2, by being initially given the same kind of knowledge the child had and by being exposed to language in the same kind of contexts as the child did. The program learned language according to a reasonable progression, making the same sort of errors that children do at intermediate stages. No parts of speech or traditional grammatical constructions are learned. It also acquires structural knowledge after knowledge of meaning, because no structural knowledge can be associated with a word until the meaning of that word is learned. This This aspect of the model offers an explanation for why children learn structure following meaning (Wetstone and Friedlander, 1973). In addition to English, the program has been tested on comparable subsets of Japanese, Russian, Chinese, Hebrew, and Spanish. Its performance with these languages was equivalent to its learning of English, suggesting that the program has no English-specific mechanisms.

This research suggests several conclusions. It suggests that a large part of the language learning problem lies in accounting for how the child infers the meaning of the language he hears. It argues that the mechanisms underlying the learning of meaning and structure are the same. It questions the role of traditional grammatical models both in language learning and language understanding, and suggests that models of language learning must be based on strong models of language understanding. In particular, it questions Chomsky's (1980) position that language is not learned. This work suggests that plausible learning models of language development are possible.

Further research should proceed in many directions. In particular, the program discussed here should be extended to model the development of comprehension of more complex constructions, such as relative clauses, and the generation of language.

Acknowledgements

Dr. Roger Schank's assistance in this work was invaluable. Peter Selfridge provided useful comments on this paper.

Bibliography

Birnbaum, L., and Selfridge, M. (1979). Problems in Conceptual Analysis of Natural Language. Research Report 168, Department of Computer Science, Yale University.

Bruner, J.S., Goodnow, J. J., and Austin, G.A., (1956). A Study of Thinking. John Wiley and Sons, New York

Chomsky, N., (1980). Rules and Representations, excerpted from Rules and Representations. Columbia University Press, New York

Hoodenraad, R., Grieve, R., Baldwin, P., and Campbell, R. (1976). Comprehension as an Interactive Process. In R. N. Campell and P. T. Smith (eds.) Recent Advances in the Psychology of Language., Plenum Press, New York

Newport, E.L., (1976). Motherese: the Speech of Mothers to Young Children. in N.J. Castellan, D.B. Pisoni, and G.R. Potts, (eds.) Cognitive Theory: VI II., Lawrence Erlbaum Assoc., Hilsdale, N.J.

Schank, R. C., (1973). Identification of Conceptualizations Underlying Natural Language. In R. C. Schank and K. M. Colby (eds.) Computer Models of Thought and Language W.H. Freeman and Co., San Fransisco.

Schank, R. C., and Selfridge, M. (1977). How to Learn/What to Learn. in Proceedings of the International Joint Conference on Artificial Intelligence, Cambridge, Mass.

Selfridge, M. (1980). A Process Model of Language Acquisition. Computer Science Technical Report 172, Yale University, New Haven, Ct.

Strohrer, H. and Nelson, K.E., (1974). The Young Child's Development of Sentence Comprehension: Influence of Event Probability, Non-verbal Context, Syntactic Form, and Strategies. Child Dev., 45:567-576

Westone, H. and Friedlander, (1973). The Effect of Word Order on Young Children's Responses to Simple Questions and Commands. Child Dev. 44:734-740