

A COMPUTER MODEL OF CHILD LANGUAGE ACQUISITION

Mallory Selfridge

Department of EE and CS
University of Connecticut
Storrs, CT. 06268

ABSTRACT

Children learn different aspects of language in a characteristic order, and make characteristic errors during acquisition. This paper focusses on five specific data, and explains these data in terms of two hypotheses regarding the relationship of comprehension to generation and the relationship between meaning and syntax. A computer model, CHILD, is described which embodies these hypotheses and which manifests the same five data. CHILD's performance suggests that the explanations are plausible.

I INTRODUCTION

The question of why children's language acquisition has the distinct characteristics it does is an important question for the development of computer models of language learning. This paper considers five specific characteristics of child language learning. It proposes explanations for them based on a computer model of language learning, CHILD, which displays those characteristics.

These characteristics involve the order of acquisition of various aspects of language, and errors produced during development. First, comprehension of certain aspects of language precedes the generation of those aspects [1]. Second, generation is thought to develop from shorter utterances to longer utterances [3]. Third, knowledge of language meaning precedes knowledge of language syntax [12]. Fourth, children make comprehension errors when the intended meaning of an utterance differs from the semantically most likely meaning. [8,5]. Fifth, children misunderstand utterances whose syntax indicates an interpretation different from that which might be obtained using a semantics-only interpretation. In such cases, younger children will respond as if they produced the semantics-only interpretation, while older children will respond correctly using syntactic knowledge [11].

In order to address these issues, this paper explores two hypotheses: that knowledge of language syntax is indexed under knowledge of language meaning, and that learning to generate language is driven by learning to understand language. These hypotheses imply that what children learn during language development is the contents of their lexicon, and that the processes of comprehension, generation, and inference are not learned.

The first hypothesis, the "semantically-indexed syntax hypothesis", is focussed around the representation of word meaning as conceptual case frames. The process of understanding is seen as the problem of inferring the filler for a case slot in the meaning of one word from among the meanings of the other words in the utterance. Syntax functions in situations in which semantic information is insufficient, to supply positional information as to where in the utterance a filler might be found. The process of generation is seen as finding a word to express the top level case frame of a concept, and finding other words to express case fillers not spanned by the meaning of the first word. The purpose of syntax in generation is to specify the appropriate position of the words to express the fillers relative to the word expressing the case frame, and relative to words expressing other fillers. The semantically-indexed syntax hypothesis assumes that comprehension and generation both access the same word meanings and syntactic knowledge. It states that syntactic knowledge is stored associated with a slot in a case frame which is the meaning of a word.

The second hypothesis, the "comprehension-driven generation hypothesis", assumes that the principle effect of learning to understand is the addition to the lexicon of the meanings of new words, the refinement of the meanings of words already in the lexicon, the addition of syntactic knowledge to the meanings of words in the lexicon, and the refinement of that syntactic knowledge. It states that the ability to generate language grows as the lexicon develops during the development of comprehension.

These two hypotheses are certainly not the only factors in language development, but they may help explain certain aspects of acquisition. In particular, a computer model of language acquisition which incorporates them manifests the same five characteristics of child language development described above.

II THE CHILD PROGRAM

CHILD is a computer model of the development of language comprehension and generation abilities in a child, Joshua, between the ages of one and two years. CHILD 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 also learns to generate descriptions of the actions of others it "observes" which are appropriate to Joshua's age 2 level. CHILD learns by being given the kind of language experiences Joshua had between the ages of 1 and 2, and making use of rules to a) infer the meaning of utterances, b) attend to words, and c) learn language meaning and syntax. While CHILD only comprehends imperatives, and only generates declaratives, these are a significant subset of Joshua's language abilities at age 2, and thus are appropriate subjects for computer modelling. Further, CHILD focusses only on modelling how Joshua learned the words for concepts which he already knew; since at age 1 his conceptual repertoire was large, this appears reasonable.

The first component of a computer model of the development of Joshua's comprehension is Joshua's knowledge prior to his language learning. Observations of Joshua, reported in [10], suggest that Joshua had considerable knowledge of objects, actions, spatial relations, and gestures at age 1. CHILD contains such knowledge represented in Conceptual Dependency (CD) [9].

CHILD models Joshua's knowledge of word meaning and syntax by maintaining a "dictionary" containing words and their associated meaning and syntax. Meanings are represented as CD structures, while syntactic knowledge is represented using "sequential structure" [10] which specifies the location of the fillers for slots in CD concepts in CHILD's short-term memory.

The language understanding process which CHILD uses is based on an adult-level language understanding program called CA [2]. CA operates by reading the words of a sentence one by one, left to right. When it reads a word, it places the CD concept representing its meaning on the CLIST, an immediate processing memory. Then, the semantic and structural requirements associated with the slots in the concepts from previous words are checked to see if the new concept can be used to fill any of them. Likewise, the requirements associated with the new concept are checked to see if they can be filled by any of the concepts already on the CLIST. In this way, the meaning of the whole sentence is built up. When there are no more words, the resulting combination of concepts is made available to inferential processing, which uses more global world knowledge to combine any uncombined concepts and fill any unfilled slots.

While children say things for many reasons [7], CHILD only models one aspect of Joshua's generation ability: his ability to describe ongoing actions of others. CHILD's generation component has been developed from an adult-level generator called CGEN [6]. CGEN begins when a knowledge structure is placed on its CLIST, its immediate processing memory. When CGEN finds a concept on the front of the CLIST, it searches its dictionary to find a word whose CD meaning matches the concept being generated. When it finds one, it places this word on the front of the CLIST. When CGEN finds a word on the front of the CLIST, CGEN "says" the word by printing it. If, however, the

dictionary contains only a word whose meaning matches the root CD concept of the overall concept being expressed, then CGEN remembers that word on the CLIST, and inserts the CD concepts which filled slots in the root concept into their appropriate places on the CLIST. The word which matched the root concept may have supplied structural knowledge to determine where in the CLIST the fill-concepts should be inserted. If not, their position is arbitrary. If CGEN fails to find a word whose meaning even partially matches the concept to be expressed, it generates "um". CGEN continues by examining the front of the CLIST, either printing the words it finds, or looking up words for the concepts. Thus, when CHILD wants to generate a concept, it places that concept on CGEN's CLIST, and CGEN proceeds to express that concept in English as well as it can.

In order to learn, CHILD must infer the meaning of utterances it hears. Since at the age of 1 Joshua had knowledge of the function and properties of objects, knowledge of the meanings of gestures, and knowledge of actions, CHILD has been given this knowledge, and uses it to infer utterance meaning. If CHILD has learned some of the words in the utterance, it can also use its partial understanding as it infers utterance meaning. Since the utterances it hears are imperative, CHILD can verify the inference by "behaving" and awaiting parental approval or disapproval. In order to model Joshua's knowledge of such parental responses, CHILD has been given knowledge of the keywords "yes" and "no", and of such gestures as head nodding, head shaking, and smiling. If none of these are provided after a behavior, CHILD assumes that implicit approval has been given.

For Joshua to learn language at his age 2 level beginning with his age 1 conceptual knowledge required the experiences he had between the ages of 1 and 2. These experiences were primarily ones in which Joshua could infer the meaning or much of the language he heard. Since Joshua received this kind of experiences, CHILD has been provided with similar experiences.

CHILD's strategies for learning word meaning and syntax are still under development; the current strategies have proven adequate for the tasks undertaken so far. The basic orientation of these learning strategies corresponds to learning rules proposed by a number of others [4,13].

CHILD takes the meaning of a word to be that component of the utterance meaning which is common to all occurrences of that word. This means that CHILD cannot learn words with multiple meanings. Further, it is assumed that the meanings of words appearing in the same utterance do not overlap. Thus if some words in the utterance are already known, then their meaning can be factored out of the meaning of the utterance as a whole and only the remainder considered as the meaning of the unknown word. Thus the first time a word is heard, it is given a meaning corresponding to the inferred meaning of the entire utterance with meanings or known words factored out. Each further time a word occurs in an utterance, the meaning of

that utterance is then "intersected" with the word's current meaning and the result is the word's new meaning. If the intersection is empty, then the word is considered meaningless and marked as such.

CHILD's syntax learning strategy require that a representation of the meaning of the utterance be obtained, along with a copy of the short-term memory used during the analysis phase. The positions of the words whose meanings fill slots in the meanings of other words are remembered, and the sequential structure of their position is stored under that slot of the meaning of the word in the word's definition.

III CHILD'S PERFORMANCE

Most of CHILD was originally developed at Yale University, while current work proceeds at the University of Connecticut. The excerpts have been taken from a complete run of CHILD, and are presented in the order in which they appeared in that. Each excerpt represents CHILD at a different stage in development as it progresses from Joshua's age 1 abilities to Joshua's age 2 abilities. Messages concerning CHILD's internal state, such as rules being activated, have been largely omitted. This run represents one possible developmental progression for CHILD. Other initial knowledge or other language experiences would lead to a different progression.

At the onset, CHILD has only the kind of knowledge ascribed to Joshua age 1. In the excerpt below, the "parent" types a lowercase utterance to CHILD, which responds with a message stating its lack of understanding. When the parent provides gestures via simulated visual input, CHILD 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 RESPONDS:
| (ATRANS ACTOR (CHILD) OBJECT (BALL1)
|   TO (POSS VAL (PARENT)))
```

Thus at the stage corresponding to Joshua at age 1, CHILD understands no language, though it will respond to gestures that it knows. At this stage its generation abilities are also poor. As the following example shows, CHILD has little to say about an action that it observes.

```
|CHILD SEES: (PARENT PUTS BALL1 ON TABLE1)
|
|INTERPRETING GESTURE AS:
| (PTRANS ACTOR (PARENT) OBJECT (BALL1)
|   TO (TOP VAL (TABLE1)))
|
|CHILD SAYS: um
```

Hence CHILD's initial stage is that in which it knows neither word meanings nor syntax. Its transition out of the first stage occurs when it begins to learn word meaning; for example the meaning of "ball." When CHILD knows the meaning of one word, it can demonstrate rudimentary understanding. In the excerpt below, the parent tells child to "give me the ball." CHILD understands only the word "ball", and must use its knowledge of possible actions that can involve balls to infer possible responses. The first response is not what the parent intended, so CHILD tries again, successfully.

```
|PARENT SAYS: give me the ball
|
|CHILD RESPONDS:
| (PTRANS ACTOR (CHILD) OBJECT (BALL1)
|   TO (TOP VAL (TABLE1)))
|
|PARENT SAYS: no
|
|CHILD RESPONDS:
| (ATRANS ACTOR (CHILD) OBJECT (BALL1)
|   TO (POSS VAL (PARENT)))
|
|PARENT SAYS: (smiles and nods)
```

CHILD at this stage can respond to words and generate responses on the basis of its understanding of them, as shown in the following.

```
|CHILD SEES: (PARENT PUTS BALL1 ON TABLE1)
|
|INTERPRETING GESTURE AS:
| (PTRANS ACTOR (PARENT) OBJECT (BALL1)
|   TO (TOP VAL (TABLE1)))
|
|CHILD SAYS: ball
```

In the third stage, shown in the excerpts below, CHILD has learned the meaning of several words, and understands some utterances correctly. In this case, it has learned the words "put", "ball", and "box". Notice, however, that although it responds correctly to the first utterance given by the parent, it misunderstands the second. This sort of error is reported in [8]. Not knowing "on", CHILD incorrectly infers that the appropriate relationship is containment, based on its knowledge of the usual relationships that obtain between balls and boxes.

```
|PARENT SAYS: put the ball in the box
|
|CHILD RESPONDS:
| (PTRANS ACTOR (CHILD) OBJECT (BALL1)
|   TO (CONT VAL (BOX1)))
|
|PARENT SAYS: put the ball on the box
|
|CHILD RESPONDS:
| (PTRANS ACTOR (CHILD) OBJECT (BALL1)
|   TO (CONT VAL (BOX1)))
```

CHILD can also use its increased knowledge of word meaning to generate better descriptions of observed actions, as shown in the following:

|CHILD SEES: (PARENT PUTS BALL1 IN BOX1)

|INTERPRETING GESTURE AS:

| (PTRANS ACTOR (PARENT) OBJECT (BALL1)
| TO (CONT VAL (BOX1)))

|CHILD SAYS: ball put

CHILD's generative ability increases with the learning of additional word meaning. In this case it has learned the additional object words "slippers", "piano", "ball", "block", and "table." CHILD can use these new words to generate better descriptions of parental action.

|CHILD SEES: (PARENT PUTS SLIPPERS1 IN BOX1)

|INTERPRETING GESTURE AS:

| (PTRANS ACTOR (PARENT) OBJECT (SLIPPERS1)
| TO (CONT VAL (BOX1)))

|CHILD SAYS: box slippers put

Although CHILD has learned enough words to express certain concepts at the three word stage, it knows no syntax. It is now taught syntax for the word "put", and learns that the fillers of the OBJECT and TO slots are found FOLLOWING "put", and that the TO slot is found FOLLOWING the OBJECT filler. This new syntactic knowledge enables CHILD to generate the description of the observed action with the words in the correct order:

|CHILD SEES: (PARENT PUTS SLIPPERS1 IN BOX1)

|INTERPRETING GESTURE AS:

| (PTRANS ACTOR (PARENT) OBJECT (SLIPPERS1)
| TO (CONT VAL (BOX1)))

|CHILD SAYS: put slippers box

The transition to the next stage occurs when CHILD learns the meaning of "on." This allows it to understand the unlikely utterance, "Put the slippers on the piano", whereas at any earlier stage it would not have understood this correctly. This is shown in the following example.

|PARENT SAYS: put the slippers on the piano

|CHILD RESPONDS:

| (PTRANS ACTOR (CHILD) OBJECT (SLIPPERS1)
| TO (TOP VAL (PIANO1)))

Although this stage represents Joshua's age 2 understanding of word meaning, CHILD has not yet learned sufficient syntax to interpret utterances whose meanings contradict the syntactic cues. It interprets the following unlikely utterance incorrectly in accord with its knowledge of the usual relationships between objects. This sort of error is similar to that reported in [11].

|PARENT SAYS: put the table on the ball

|CHILD RESPONDS:

| (PTRANS ACTOR (CHILD) OBJECT (BALL1)
| TO (TOP VAL (TABLE1)))

CHILD's generation at this stage also reflects the lack of syntactic knowledge. It knows the syntax for "put," but not for "on." This leads it to generate utterances whose order is not totally correct.

|CHILD SEES: (PARENT PUTS BOX1 ON TABLE1)

|INTERPRETING GESTURE AS:

| (PTRANS ACTOR (PARENT) OBJECT (BOX1)
| TO (TOP VAL (TABLE1)))

|CHILD SAYS: put box table on

The progression to CHILD's final stage is accomplished by learning the syntax of "on". When CHILD learns that the filler of the VAL slot of "on's" meaning is found following "on", it can then understand utterances whose syntax contradicts the most likely meaning of the utterance. CHILD is asked to "Put the table on the ball," and its correct response to this utterance should be interpreted as modelling a child who understands and who might go over to the table and tug at it a bit while looking at the ball.

|PARENT SAYS: put the table on the ball

|CHILD RESPONDS:

| (PTRANS ACTOR (CHILD) OBJECT (TABLE1)
| TO (TOP VAL (BALL1)))

At this stage, CHILD also finally generates a description of the observed action correctly.

|CHILD SEES: (PARENT PUTS SLIPPERS1 ON TABLE1)

|INTERPRETING GESTURE AS:

| (PTRANS ACTOR (PARENT) OBJECT (SLIPPERS1)
| TO (TOP VAL (TABLE1)))

|CHILD SAYS: put slippers on table

This progression shows that CHILD's performance manifests those language characteristics described earlier. First, its comprehension of language meaning and syntax precedes its generation of that same meaning and syntax. Second, the length of its utterances increases during the progression. Third, its knowledge of language meaning precedes its knowledge of language syntax, since it comprehends and generates utterances without syntax before it comprehends and generates them with syntax. Fourth, CHILD incorrectly understands utterances on the basis of the semantically most likely meaning when it doesn't know all of the words in the utterance. Fifth, it misunderstands utterances whose syntax indicates an interpretation different from the semantically most likely interpretation before it has learned the syntax for the words in the utterance.

IV CONCLUSIONS

CHILD's manifestation of the five data supports both the semantically-indexed syntax hypothesis and the comprehension-driven generation hypothesis. In particular, it offers explanations

for this data. Generation follows comprehension because generation cannot occur until comprehension learning adds words to the dictionary. Length of utterance increases because the number of words available to express a concept increases during comprehension learning. Knowledge of language meaning precedes knowledge of language syntax because syntax is indexed under word meanings and hence cannot be learned before the word meaning. Misunderstanding of utterances whose correct interpretation is not the semantically most probable occurs because children use their knowledge of probable meanings to augment gaps in their understanding. Finally, misunderstanding of utterances whose syntax suggests an interpretation different from the semantically most likely interpretation occurs because knowledge of syntax is learned after knowledge of meaning. This error is made when enough knowledge of meaning has been acquired to produce an interpretation, but not enough syntax has been learned to produce a correct interpretation.

There are certainly many other factors in child language acquisition which have not been considered here, but this paper offers support for the semantically-indexed syntax hypothesis and the comprehension-driven generation hypothesis as components of a complete model of child language acquisition. Further research must address questions posed by the data of more advanced child language acquisition, which could support other methods of indexing syntactic knowledge and other possible relationships between comprehension and generation.

ACKNOWLEDGEMENTS

Dr. Roger Schank's assistance in this work was invaluable. Dr. Richard Cullingford has made many valuable suggestions, and Dr. Katherine Nelson has provided many insights into problems of modelling child language learning. Peter Selfridge, Larry Birnbaum, Marie Bienkowski and Jamie Callan have contributed both ideas and criticisms.

REFERENCES

- [1] Benedict, H. (1976). Language Comprehension in 10 to 16 Month-old Infants. (thesis) Department of Psychology, Yale University, New Haven, Connecticut.
- [2] Birnbaum, L., and Selfridge, M. (1981). in Inside Artificial Intelligence: Five Programs plus Miniatures. Schank R. and Riesbeck C.K. (eds.), Lawrence Erlbaum Associates, Hillsdale, NJ.
- [3] Bloom, L., 1973. One Word at a Time. Mouton, The Hague.
- [4] Bruner, J.S., Goodnow, J. J., and Austin, G.A., (1956). A Study of Thinking. John Wiley and Sons, New York
- [5] Clark, E.V. and Clark, H.H., (1978). Psychology and Language, Harcourt Brace Jovanovich.

- [6] Cullingford, R.E., Krueger, M.W., and Selfridge, M. (1981) A Picture and a Thousand Words: Automated Explanations as a Component of a Computer-Aided Design System. Forthcoming.
- [7] Halliday, M.A.K., (1975). Learning How to Mean: Explorations in the Development of Language Edward Arnold, London.
- [8] Hoodenraad, R., Grieve, R., Baldwin, P., and Campbell, R. (1976). Comprehension as an Interactive Process. In R. N. Campbell and P. T. Smith (eds.) Recent Advances in the Psychology of Language, Plenum Press, New York
- [9] 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 Francisco.
- [10] Selfridge, M. (1980). A Process Model of Language Acquisition. Computer Science Technical Report 172, Yale University, New Haven, Ct.
- [11] 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
- [12] 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
- [13] Winston, P. (1975). Learning Structural Descriptions from Examples. In P.H. Winston (ed.) The Psychology of Computer Vision. McGraw-Hill, New York.