

How do Children Learn to Judge Grammaticality?
A Psychologically Plausible Computer Model

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

1.0 Introduction

If a young child is asked whether the sentence "ball me the throw" sounds "silly" or "ok", chances are the child will respond "silly." Encouraged to "fix it up," the child may well generate "throw me the ball." Such behavior was reported by Gleitman et al. [7] for children of two-and-a-half and five years. It implies that by these ages children have acquired at least some ability to judge a sentence's grammaticality. Further, Gleitman et al. report that by age five, children's judgements increase in sophistication. Thus children's ability to judge grammaticality apparently increases as they learn language.

Unfortunately, little is known about the mechanisms responsible for the development of such abilities. Pinker [10], reviewing language acquisition models, reports no work in this direction. Anderson's [1] model of language learning does not address learning to make grammaticality judgements. Recent research (e.g. [2,3,8]) on syntactic recognition and learning has not been integrated into a model of child learning. The question remains: "how do children learn to make grammaticality judgements?"

This paper addresses this question by proposing a three stage model, implemented and tested in the CHILD program [12,13,14,15]. During stage one CHILD knows word meanings but not syntax, and can understand sentences, but cannot tell that word order is incorrect. During the second stage, CHILD has learned active syntax, and notices incorrect word order for active sentences. During the third stage, CHILD learns passive syntax, and notices incorrect word order for both active and passive sentences. This progression corresponds generally to Gleitman et al.'s finding that as children learn more language their ability to make grammaticality judgements increases.

CHILD's mechanisms may provide part of the answer to the problem of how children learn to make grammaticality judgements of sentences with incorrect word order. These mechanisms have been developed to account for a number of different data about child language learning [14], and their extension to the problem of grammaticality judgements has been straightforward. The CHILD model suggests that children learn to make such judgements almost entirely as a side-effect of mechanisms whose primary function is directed elsewhere. This paper describes the CHILD program, and presents sample output. The question of learning to make grammaticality judgements is considered, and several predictions are described which may confirm or deny this account.

2.0 The CHILD Program

CHILD is a computer model of the development of language comprehension and generation abilities written in Franz LISP and currently running on a DEC VAX 11/780. It begins with world knowledge and language experiences similar to those received by

children, and learns a subset of the word meaning and syntax which children learn. After learning, CHILD can correctly understand utterances which it previously misunderstood, and can generate English describing events it "observes."

CHILD's language comprehension process is a version of the CA program [4] which incorporates mechanisms derived from Wilks' [16] preference parsing. CHILD's analysis process combines Conceptual Dependency (CD) [11] word meanings to form a CD representing the meaning of the entire utterance. Understanding begins when the meanings of input words are placed in a short term memory. CHILD then retrieves semantic requirements associated with those slots but specific to that particular word. It searches the short term memory for a word meaning which best satisfies those features, and fills the empty slot with that meaning. The syntactic features are formed from the positional predicates PRECEDES and FOLLOWS. These relate the position of a candidate slot filler to either the word they were stored under, a filler of another slot in that word's meaning, or a lexical function word. Each slot in the meaning of a word has a collection of features describing where in the input a filler is expected to be. In order to understand different voices, these features are organized into disjunctive "feature sets." Each set characterizes one order in which slot fillers appear. During understanding, feature set selection is performed by considering which set most successfully characterizes the input.

CHILD learns syntax by acquiring syntactic features and build disjunctive feature sets. After having understood an utterance, CHILD examines a record of the input, and examines the meaning of every input word. It then examines every empty slot in each such meaning. It accesses the record of the input to find where in the input the filler for that slot occurred. It creates a description of this position using PRECEDES and FOLLOWS. CHILD must then decide whether this description constitutes a new feature set or should be merged with an existing feature set. CHILD's strategy is based on a suggestion by Iba [8]. CHILD compares the features extracted from the current input with any existing feature sets: the position description is merged with a previous set only if one is a subset of the other. Otherwise, the description is learned as a new feature set.

CHILD notices that a sentence is ungrammatical if any syntactic features within the selected feature set characterizing the position of a slot filler are not true of the position of that filler. CHILD uses these features to generate an explanation of why the sentence was ungrammatical, and uses its language generation abilities [6] to generate a correct version based on the sentence's understood meaning according to whatever word meaning and syntax it knows about at that stage of learning.

3.0 Learning to Make Grammaticality Judgements

The following example is edited from a complete

run of the program during which it learns meaning and syntax for all the words it knows. The example begins after CHILD has learned meanings for "throw", "me", "Child", "Mom", "on", "table," and "ball." For this example, the meaning of "throw" has been simplified from a complex CD into \$THROW and processing of "the" has been ignored. As shown below, CHILD is initially given an ordinary sentence which it understands correctly. The second sentence has incorrect word order. CHILD understands this sentence correctly, but fails to notice the incorrect order.

```
CHILD hears "throw me the ball"
CHILD understands
($THROW ACTOR (CHILD) OBJECT (BALL1) TO (PARENT1))

CHILD hears "ball me the throw"
CHILD understands
($THROW ACTOR (CHILD) OBJECT (BALL1) TO (PARENT1))
```

Transition to the second stage occurs when CHILD learns active syntax for "throw." This occurs when it hears an example sentence whose interpretation is unambiguous, and has heard the word a number of times without modifying its meaning. Given this sentence, CHILD notes the positions of the fillers (summarized in linear order here), and stores them in a feature set under the word "throw."

```
CHILD hears: "throw me the ball"
CHILD's understanding is:
($THROW ACTOR (CHILD) OBJECT (BALL1) TO (PARENT1))

CHILD learns syntax of "throw"
order is: $THROW TO OBJECT

ATTEMPTING MERGE OF NEW FEATURES WITH EXISTING SET
NO EXISTING FEATURES SETS
CREATING NEW FEATURE SET
```

Having learned active syntax for "throw," CHILD uses this knowledge during stage 2 understanding. It notices when the word order of a sentence is incorrect, as the first sentence below shows. CHILD prints out the reasons it thought the sentence was incorrect, and generates a correct version from the understood meaning of the sentence. However, CHILD also decides that a passive sentence with correct order is incorrect, as the second sentence below demonstrates.

```
CHILD hears "ball me the throw"
CHILD understands
($THROW ACTOR (CHILD) OBJECT (BALL1) TO (PARENT1))

INCORRECT SENTENCE NOTICED:
"throw" should precede "ball"
"throw" should precede "me"
CORRECTION: "throw me the ball"

CHILD hears
"the ball was throw n on the table by Child"
CHILD understands
($THROW ACTOR (CHILD) OBJECT (BALL1)
  TO (TOP VAL (TABLE1)))

INCORRECT SENTENCE NOTICED:
"Child should precede "throw"
"throw" should precede "ball"
CORRECTION: "Child throw ball on table"
```

The transition to the third stage occurs when CHILD

learns passive syntax for "throw," and creates a second feature set for the new syntactic features.

```
CHILD hears: "the ball was throw n to Mom by
  Child"
CHILD's understanding is:
($THROW ACTOR (CHILD) OBJECT (BALL1) TO (PARENT1))

CHILD learns syntax of "throw"
order is: OBJECT "was" $THROW "to" TO "by" ACTOR

ATTEMPTING MERGE OF NEW FEATURES WITH EXISTING SET
MERGE FAILS
CREATING NEW FEATURE SET
```

Once CHILD has learned passive syntax (reported in more detail in [15]), it can then judge passive sentences. It correctly judges the passive sentence which it previously judged incorrect. The second sentence below is an incorrect passive, and CHILD correctly understands it, prints out the reasons it was judged incorrect, and generates its corrected version.

```
CHILD hears
"the ball was throw n on the table by Child"
CHILD understands
($THROW ACTOR (CHILD) OBJECT (BALL1)
  TO (TOP VAL (TABLE1)))

CHILD hears: "the ball to Mom throw n by Child"
CHILD's understanding is:
($THROW ACTOR (CHILD) OBJECT (BALL1) TO (PARENT1))

INCORRECT SENTENCE NOTICED:
"the ball" should precede "was"
"thrown" should precede "to Mom"
CORRECTION: "the ball was thrown to Mom by Child"
```

As shown above, CHILD does progress through a series of stages which generally correspond to data reported by Gleitman et al., during which it learns to make increasingly accurate and complex grammaticality judgements. Initially knowing no syntax, all sentences are judged correct. After learning active syntax, it successfully judges active sentences, but judges passive sentences as if they should have been active. Upon learning passive syntax, CHILD judges both active and passive sentences correctly. In the complete run, CHILD also learns to understand noun phrases, prepositional phrases, and adverbial phrases, and learns to make judgements about sentences containing these constructions.

4.0 How Do Children Learn to Make Grammaticality Judgements?

CHILD's answer to this question depends upon a number of factors: a) the representation of language syntax as a set of independent features characterizing the position in the input where a slot filler may occur; b) learning of syntactic features while learning to understand; c) the evaluation of syntactic features and semantic preferences as a necessary part of understanding. Given these mechanisms, children make grammaticality judgements by analyzing syntactic violations occurring during understanding. They generate correct versions of incorrect sentences by applying their language generation ability to the understood meaning of that sentence. Thus children acquire the ability to make grammaticality judgements as a side effect of acquiring syntactic features needed for understanding.

This account of learning to make grammaticality judgements makes several predictions. First, this model predicts that people's judgements of incorrect sentences will not merely be "grammatical" or "not grammatical," but rather judgements as to the relative grammaticality of a sentence. This prediction follows from CHILD's generation of a number of different reasons for a sentence's incorrectness. Second, this model predicts that as a child learns increasing amounts of syntax he will find certain sentences increasingly ungrammatical. This is because newly learned syntax becomes available to judge grammaticality, and thus the number of violated syntactic features increases. Third, this model predicts that before learning passive syntax children will judge non-reversible passive sentences to be ungrammatical. This is because at this stage they are using active syntax to understand passive sentences. Later, when they have learned passive syntax, they will no longer judge non-reversible passive sentences ungrammatical.

Clearly, this work has not completely solved the problem of how children learn to make grammaticality judgements, since there are certainly a large number of complex syntactic constructions which CHILD cannot handle. In addition, it is not even clear what exactly constitutes such judgements, since Gleitman et al. report that children think sentences are "silly" for a number of reasons not discussed here. It is hoped, however, that this approach will prove a promising direction for further research, since it is grounded in mechanisms which manifest and explain a number of other psychological data.

Acknowledgements

Thanks to Rich Cullingford for sponsoring this paper, and to Peter Selfridge, Oliver Selfridge, Don Dickerson, Jason Engleberg, and Marie Bienkowski for helpful discussions of this work and for commenting on drafts of this paper.

References

- [1] Anderson, J.R. (1981). A Theory of Language Acquisition Based On General Learning Principles. Proc. 7th IJCAI, Vancouver, Canada.
- [2] Baker, C.L. and McCarthy, J.J. (1981). The Logical problem of Language Acquisition. M.I.T. Press, Cambridge, Mass.
- [3] Berwick, R.C. (1977). Learning Structural Descriptions of Grammar Rules from Examples. Proc. 5th IJCAI, Cambridge, Mass.
- [4] Birnbaum, L., and Selfridge, M. (1981). Conceptual Analysis of Natural Language, in Inside Computer Understanding: Five Programs plus Miniatures. Schank R. and Riesbeck C.K. (eds.), Lawrence Erlbaum Associates, Hillsdale, NJ.
- [5] Chomsky, N. (1965). Aspects of the Theory of Syntax. M.I.T. Press, Cambridge, Mass.
- [6] Cullingford, R.E., Krueger, M.W., Selfridge, M. and Bienkowski, M. (1981). Automated Explanations as a Component of a CAD System. IEEE Trans. SM&C. Special Issue on Human Factors and User Assistance in CAD, December, 1981.
- [7] Gleitman, L.R., Gleitman, H. and Shipley, E.F. (1972). The Emergence of the Child as Grammarian, Cognition, 1-2/3:1-164.
- [8] Iba, G. (1979). Learning Disjunctive Concepts from Examples. M.I.T. A.I. Memo #548, M.I.T., Cambridge, Mass.
- [9] Marcus, M. (1980). A Theory of Syntactic Recognition for Natural Language. M.I.T. Press, Cambridge, Mass.
- [10] Pinker, S. (1979). Formal Models of Language Learning. Cognition, 7:217-283.
- [11] 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.
- [12] Selfridge, M. (1980). A Process Model of Language Acquisition. Computer Science Technical Report 172, Yale University, New Haven, Ct.
- [13] Selfridge, M. (1981a). Why Do Children Say "Goed"? A Computer Model of Child Generation. Proc. Third Annual Meeting of the Cognitive Science Society. Berkeley, CA.
- [14] Selfridge, M. (1981b). A Computer Model of Child Language Acquisition. Proc. 7th Int. Joint Conf. on Artificial Intelligence. Vancouver, Canada.
- [15] Selfridge, M. (1982). Why Do Children Misunderstand Reversible Passives? The CHILD Program Learns to Understand Passive Sentences. Submitted to the 3rd Annual AAAI Conference, Pittsburgh, Penn.
- [16] Wilks, Y., (1973). Parsing English II. In E. Charniak and Y. Wilks (eds.) Computational Semantics. North-Holland Publishing Co., NY, NY.