

# Integrated Learning of Words and their Underlying Concepts\*

Raymond J. Mooney

Coordinated Science Laboratory  
University of Illinois at Urbana-Champaign

## Abstract

Models of learning word meanings have generally assumed prior knowledge of the concepts to which the words refer. However, novel natural language text or discourse can often present both unknown concepts and words which refer to these concepts. Also, developmental data suggests that the learning of words and their concepts frequently occurs concurrently instead of concept learning proceeding word learning. This paper presents an integrated computational model for acquiring both word meanings and their underlying concepts concurrently. This model is implemented as a word learning component added to the GENESIS explanation-based schema acquisition system for narrative understanding. A detailed example is described in which GENESIS learns provisional definitions for the words "kidnap", "kidnapper", and "ransom" as well as a kidnapping schema from a single narrative.

## Introduction

Previous computational models of the acquisition of word meaning [Berwick83, Granger77, Selfridge82] have assumed existing knowledge of the concept underlying the word to be learned. In these models, word learning is a process of using surrounding context to establish an identification between a new lexical item and a known concept. However, new words are not always encountered as labels for known concepts. When encountering a new concept in natural language text or discourse, it is quite likely that one will also come across unknown words which refer to various aspects of the new concept. A word learning model which requires prior knowledge of the underlying concept will be unable to acquire even provisional meanings for such words.

Developmental studies suggest that the learning of words and their underlying concepts frequently occurs concurrently. Experiments by Gopnik and Meltzoff revealed that children's acquisition of "disappearance" words occurred at about the same time they learned to solve object-permanence tasks involving invisible displacements [Gopnik86]. From this data, they concluded that learning may often involve "concurrent cognitive and semantic developments, rather than involving cognitive prerequisites for semantic developments." Bowerman [Bowerman80] and Kuczaj [Kuczaj82] have also used developmental data to argue for an interactive approach to language and concept acquisition.

This paper describes an integrated computational model of the acquisition of word meanings and their underlying concepts. This approach was developed in an attempt to add word learning abilities to the GENESIS explanation-based schema acquisition system [Mooney85]. From a single natural language narrative, the current GENESIS system is able to acquire a new schema as well as provisional meanings for several schema related words.

## An Overview of Schema Acquisition in GENESIS

GENESIS [Mooney85] is an explanation-based learning system [DeJong86, Mitchell86] which learns a plan schema from a single instance by determining *why* a particular sequence of actions observed in a specific narrative allowed the actors to achieve their goals. During the understanding process, GENESIS attempts to construct explanations for characters' actions in terms of the goals their actions were meant to achieve. This process involves plan and script-based understanding

---

\* This research was supported by the Office of Naval Research under grant N-00014-86 K-0309.

mechanisms like those employed by previous narrative processing systems [Schank81]. When the system observes that a character has achieved an interesting goal in a novel way, it generalizes the composition of actions the character used to achieve this goal into a new schema. The generalization process (described in [Mooney86]) consists of an analysis of the causal model of the narrative which removes unnecessary details while maintaining the validity of the explanation. The resulting generalized set of actions is then stored as a new schema and used by the system to correctly process narratives which were previously beyond its capabilities.

In [Ahn87], experimental evidence is presented which indicates that, like GENESIS, people can also acquire a schema by generalizing the explanation of a single narrative. After reading one specific instance of a novel plan, subjects were able to describe the underlying schema in abstract terms, generate a different instance, and correctly answer questions about the general schema.

### An Example of Learning Words with their Concepts

This section presents an example of GENESIS' ability to learn word meanings as well as their corresponding concepts from a single example. First the system is given the following kidnapping story. At this point, GENESIS has schemata for threatening, capturing, making bargains, and a number of other actions as well as definitions for many words; however, it does not have a schema for kidnapping-for-ransom nor definitions for the words "kidnap", "kidnapper", or "ransom".

#### Story 1

Fred is Mary's father and is a millionaire. John approached Mary and pointed a gun at her. She was wearing blue jeans. He told her if she did not get in his car then he would shoot her. He drove her to his hotel and locked her in his room. John called Fred and told him John was holding Mary captive. John told Fred if Fred gave him \$250000 at Trenos then John would release Mary. Fred paid him the ransom and the kidnapper released Mary. Valerie is Fred's wife and he told her that someone had kidnapped Mary.

From this single instance, GENESIS learns a general schema for kidnapping-for-ransom (which it calls CaptureBargain based on the two main actions which compose it) as well as preliminary definitions for "kidnapper", "ransom", and "kidnap". A paraphrase of the CaptureBargain schema is shown below:

#### CaptureBargain(?x97,?a52,?b11,?c4,?y15,?i19)

?b11 is a person. ?c4 is a location. ?x97 is a character. ?b11 is free. ?x97 captures ?b11. ?a52 is a character. ?x97 contacts ?a52 and tells it that ?b11 is ?x97's captive. ?y15 is a valuable. ?x97 wants to have ?y15 more than it wants ?b11 to be ?x97's captive. ?a52 has a positive relationship with ?b11. ?a52 has ?y15. ?x97 and ?a52 carry out a bargain in which ?x97 releases ?b11 and ?a52 gives ?x97 ?y15 at ?i19.

The provisional definition learned for "kidnap" is an action describing an instance of CaptureBargain where the subject is the actor (?x97) and the direct object is the person he captures and then releases in exchange for payment (?b11). The definition conjectured for "kidnapper" is a person filling the actor role of the new schema (?x97) and the definition for "ransom" is a valuable item given to this actor (?y15). These definitions do not exactly match the standard dictionary definitions of these words, but they are reasonable approximations given their use in this one example. The lexical and schematic knowledge acquired from this example enables the system to subsequently explain the following stories as instances of its new CaptureBargain schema.

#### Story 2

Ted is Alice's husband. A kidnapper took Alice into a room. Bob got \$75000 and released Alice.

#### Story 3

Ted is Alice's husband. John took Alice into a room. Ted paid John the ransom and John released Alice.

#### Story 4

Steve kidnapped Valerie. Mike was Valerie's father and paid Steve \$30000.

Prior to learning, GENESIS could not construct explanations for any of these stories since each one requires both knowledge of the schema and a definition for the appropriate kidnap-related word.

### How GENESIS Learns Word Meanings

Since the processes for learning *role labels* such as "kidnapper" and "ransom" and that for learning *schema labels* such as "kidnap" are somewhat different, each of these procedures will be discussed separately.

#### *Learning Role Labels*

The procedure used to learn role labels is similar to the technique used by the FOUL-UP system [Granger77] except that it is integrated with the schema learning and schema activation processes. When the parser (a modified version of McDYPAR [Dyer83]) encounters an unknown word when there is an expectation for a noun of some class, a dummy variable is created, annotated with the unknown word, and allowed to fill the expectation. For example, the phrase "Fred paid him the ransom" in Story 1 is parsed into the assertion:  $Atrans(Person1, ?x1, Person3)$  where  $?x1$  is marked with the fact that it came from the unknown word "ransom." If an input pattern with an unknown-word variable like  $?x1$  matches a pattern expected by currently active schema, then a provisional definition for the word is made based on the constraints on the schema role which matches the variable. For example, in the "ransom" case, the previous sentence in Story 1 suggested a Bargain schema between John and Fred and set up expectations for the two proposed actions. Since "Fred paid him the ransom" matches the expected action "Fred gives John the \$250000," and "ransom" fills the role of the item whose possession is transferred, an initial definition is made for "ransom" stating that it is a physical object whose possession is transferred during a Bargain.

However, this is not the final definition created for "ransom" since an additional process is performed when a new schema is learned. Each of the sub-actions composing a new schema is checked for roles which are filled by unknown-word variables or which were previously matched to such a variable resulting in an initial definition. In either case, a new definition is created for the unknown word based on the role it fills in the learned schema and the schema constraints on this role. Consequently, when the CaptureBargain schema is subsequently recognized and learned from Story 1, it causes "ransom" to be redefined as a valuable item whose possession is transferred to the actor in a CaptureBargain schema ( $?y15$ ). The rationale for having *learned* schemata take precedence when defining such words is that a learned schema represents a new situation and therefore new words are assumed more likely to be directly associated with it than with an existing schema like Bargain.

The provisional definition for a role label like "ransom" contains two parts. The first is a set of constraints on the object itself, such as an assertion that it be Valuable. The second part is a suggestion of the schema of which it is a role. The fact that a role label definition can suggest a relevant schema allows GENESIS to use the definition to construct explanations for narratives which it otherwise would not understand. When the word "ransom" is subsequently encountered in Story 3, it suggests that the CaptureBargain schema might be relevant. This schema is then used in a top-down fashion to construct an explanation for the text. Since no other piece of information suggests CaptureBargain, the learned definition for "ransom" is crucial in understanding this story.

However, most words are not role labels like "kidnapper" and "ransom." For example, consider replacing the word "ransom" in Story 1 with the word "moolah." Since the word "moolah" is unknown, GENESIS gives it a definition identical to the one it learned for "ransom." In order to be able to recover from such mistakes, the system monitors the schemata suggested by newly learned words. If a new word subsequently suggests a schema which does not explain any future inputs, the suggestion is removed. Consequently, after receiving a murder-for-inheritance story in which the word "moolah" is used, "moolah" ceases to suggest CaptureBargain.

### Learning Schema Labels

Learning meanings for verbs which refer to entire plan schemata is a more difficult task since the relevant context is potentially much broader. A sentence such as "John robbed the store" may be used to introduce a long piece of text elaborating the situation, to succinctly summarize a previous piece of text, or to simply refer to a single action in an even larger plan. A number of heuristics have been developed which allow a reasonable guess to be made regarding the reference of such unknown verbs. The following one is used to resolve the meaning of "kidnap" as used in Story 1.

If one character informs another that some action occurred and a schema whose actor is the same as this action's was recently acquired from the narrative, and this schema also has roles filled by the speaker and any direct and indirect objects of the action, then assume the speaker is summarizing the event and that the unknown act refers to the new schema.

Specifically, since Fred tells his wife that "someone kidnapped Mary" and both him and Mary were participants in the just completed CaptureBargain schema, GENESIS assumes "kidnap" refers to CaptureBargain (i.e. one can summarize the schema by saying: ?x97 kidnapped ?b11). An appropriate definition is then created for "kidnap" and can be used to help understand Story 4. In this narrative, "Steve kidnapped Valerie" is interpreted as describing an instance of CaptureBargain in which Steve is the actor and Valerie is the victim. The assertions that "Mike is Valerie's father" and "Mike paid Steve \$30000" are then understood as parts of the expansion of this instance of CaptureBargain. Another heuristic involves the mentioning of an unknown action followed by an elaboration which describes a novel schema.

### Conclusions

Unlike previous approaches to the acquisition of word meanings, the present approach does not assume prior knowledge of the underlying concepts. Learning definitions for new words is integrated with an explanation-based concept learning mechanism. This allows the system to learn concepts and word meanings concurrently, which is a phenomenon which has been observed in developmental studies. However, this paper reports only preliminary work in the area of integrated word learning. There are many problems which still need to be addressed. A few of these are listed below.

- (1) The procedure for removing schema suggestions from new definitions is too strict. One counter-example should not eliminate a suggestion and repeated usefulness of a suggestion should make it resistant to elimination.
- (2) Morphology of unknown words should be considered. A "kidnapper" is clearly the actor of a "kidnapping."
- (3) More and better heuristics are needed for determining whether a word might be a *schema label* and to what schema it might refer.
- (4) Only *role labels* and *schema labels* are considered. Many words do not fall into either of these two categories.
- (5) Only integration with explanation-based learning is considered. Integration with *similarity-based learning* [Dietterich83] should also be examined.

Nevertheless, the current work demonstrates the feasibility and usefulness of integrating the acquisition of word meanings and concepts. Further research is needed in both AI and psychology to explore the potential symbiotic relationship between language and concept acquisition.

### Acknowledgements

The author would like to thank Gerald DeJong for directing this research and Dedre Gentner for pointers to the relevant developmental literature.

References

- [Ahn87] W. Ahn, R. J. Mooney, W. F. Brewer and G. F. DeJong, "Schema Acquisition from One Example: Psychological Evidence for Explanation-Based Learning." Technical Report, AI Research Group, Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, February, 1987.
- [Berwick83] R. C. Berwick, "Learning Word Meanings from Examples," *Proceedings of the Eighth International Joint Conference on Artificial Intelligence*, Karlsruhe, West Germany, August 1983, pp. 459-461.
- [Bowerman80] M. Bowerman, "The Structure and Origin of Semantic Categories in the Language Learning Child," in *Symbols as Sense: New Approaches to the Analysis of Meaning*, M. L. Foster & S. H. Brandes (ed.), Academic Press, New York, 1980, pp. 277-299.
- [DeJong86] G. F. DeJong and R. J. Mooney, "Explanation-Based Learning: An Alternative View," *Machine Learning 1, 2* (April 1986), pp. 145-176.
- [Dietterich83] T. G. Dietterich and R. S. Michalski, "A Comparative Review of Selected Methods for Learning from Examples," in *Machine Learning: An Artificial Intelligence Approach*, R. S. Michalski, J. G. Carbonell and T. M. Mitchell (ed.), Tioga Publishing Company, Palo Alto, CA, 1983, pp. 41-81.
- [Dyer83] M. J. Dyer, *In-Depth Understanding*, MIT Press, Cambridge, MA, 1983.
- [Gopnik86] A. Gopnik and A. N. Meltzoff, "Relations between Semantic and Cognitive Development in the One-Word Stage: The Specificity Hypothesis," *Child Development 57*, (1986), pp. 1040-1053.
- [Granger77] R. H. Granger, "FOUL-UP: A Program that Figures Out Meanings of Words from Context," *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, Cambridge, MA, August 1977, pp. 172-178.
- [Kuczaj82] S. A. Kuczaj, "Acquisition of Word Meaning in the Context of the Development of the Semantic System," in *Verbal Processes in Children: Progress in Cognitive Developmental Research*, C. J. Brainerd & M. Pressley (ed.), Springer-Verlag, New York, 1982, pp. 95-123.
- [Mitchell86] T. M. Mitchell, R. Keller and S. Kedar-Cabelli, "Explanation-Based Generalization: A Unifying View," *Machine Learning 1, 1* (January 1986), pp. 47-80.
- [Mooney85] R. J. Mooney and G. F. DeJong, "Learning Schemata for Natural Language Processing," *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, Los Angeles, CA, August 1985, pp. 681-687.
- [Mooney86] R. J. Mooney and S. W. Bennett, "A Domain Independent Explanation-Based Generalizer," *Proceedings of the National Conference on Artificial Intelligence*, Philadelphia, PA, August 1986, pp. 551-555. (A longer updated version appears as Technical Report UILU-ENG-86-2216, AI Research Group, Coordinated Science Laboratory, University of Illinois at Urbana-Champaign)
- [Schank81] R. C. Schank and C. Riesbeck, *Inside Computer Understanding*, Lawrence Erlbaum and Associates, Hillsdale, NJ, 1981.
- [Selfridge82] M. Selfridge, "Inference and Learning in a Computer Model of the Development of Language Comprehension in a Young Child," in *Strategies for Natural Language Processing*, W. G. Lehnert and M. H. Ringle (ed.), Lawrence Erlbaum and Associates, 1982, pp. 299-326.