Proceedings
of the
Workshop W.14
on
NATURAL LANGUAGE LEARNING

25 August 1991

12th International Joint Conference on Artificial Intelligence
(IJCAI-91)
Darling Harbour, Sydney, Australia
24 - 30 August 1991

Acquisition of Attribute Applicability

James Geller
Institute for Integrated Systems, CIS Department
and Center for Manufacturing Systems
New Jersey Institute of Technology
Newark, NJ 07102
email: geller@mars.njit.edu

ABSTRACT:

While most researchers in Machine Learning of Natural Language are interested in the acquisition of syntax, this paper focuses on the acquisition of semantics. Two techniques for the acquisition of semantics are presented, namely, the single channel and the multiple channel approach. Using the single channel paradigm, the question of how a child learns which attributes are applicable to a class of objects is addressed. A computational model for this attribute applicability is discussed in detail.

1 Introduction

In this research we are interested in modeling the acquisition of natural language semantics by computers. We are hereby guided by methods from Artificial Intelligence and Machine Learning of Natural Language [Powers89] and by the literature on child language learning. The purpose of this abstract is twofold. **(1)** We present our previous work and our general approach to semantic acquisition. **(2)** We discuss the specific problem of ***attribute applicability***.

Much of the previous work on semantic acquisition, including our own work as discussed in Section 2, has concentrated on the acquisition of concrete nouns and categories of objects. While the problems involved in these two tasks are far from resolved, we feel that it is necessary to advance the range of problems that are considered in NLML modeling.

An example of a problem that has received little attention and that appears to be the next logical step after concrete noun and category learning is the acquisition of the meaning of attributes. Like for nouns we would like to start with concrete and simple attributes that describe perceptually observable features. For example, how does a child learn what it means to be "large," "red," "round," etc., etc?

The problem of ***attribute applicability*** can be classified as being a logical precursor to the problem of attribute acquisition. It can be stated as follows:

`How does a child learn which attributes apply to a class of objects?`

This problem is an interesting research topic for a number of reasons.

1. It naturally extends the traditional work on the acquisition of prototypes and concrete nouns.

2. It does not involve all the "hard" problems of semantics, because it can be approached with a formal, logic-like approach.

3. It is a stepping stone towards modeling the acquisition of attributes.

2 Previous Work and General Approach to MLNL

Many of the publications in the area of Machine Learning of Natural Language concentrate on issues of how children and computer systems can acquire the grammar of a given language [Powers89]. There is considerably less investigation of the question how children can acquire the semantics of natural language. In previous work [Sen90,Simha90,Fung91] we have started to address the acquisition of the semantics of concrete nouns using two different techniques, the *single channel approach* and the *multiple channel approach*. Powers and Turk [Powers89] are using the terms *single modality* and *multiple modality* in a very similar sense.

In the multiple channel approach a natural language processing system receives as input pairs of noise-free pixel maps with sentences that refer to the diagrams contained in those pixel maps [Sen90]. A number of simplifying assumptions is made to permit the creation of correct associations between the diagram and the sentences. The most important one is the *novel label novel object assumption* that has been used to explain the acquisition of natural language by children [Markman90]. Another important assumption is that the sentences themselves are not ambiguous, and that a sufficient number of words are initially known to limit the number of unknown words to one per sentence. On the other hand, words that were originally unknown but became associated with an image may then be used as known words in subsequent language input.

For example, the pictorial input may look like this:

```
1111111111
 1        1
 1        1
 1        1
```

and the corresponding sentence would be:

```
This is a table.
```

Assuming that all words are known, except for "table", the system will associate the above drawing with the word "table". If we now present the diagram-sentence pair

```
  1111
1111111111
 1        1
 1        1
 1        1

The pencil is on the table.
```

the system will subtract the image of the (now known!) table and assign the only unknown word (pencil) to the image:

```
  1111
```

A paradigm similar to the one that we are using under the name multiple channel approach has recently become popular in cognitive science as a good example for a task that requires the integration of several disciplines which often tend to coexist without true integration [Feldman90]. In summary, the multiple channel approach is a simulation of early child language acquisition based on cognitive limitations of the child and ostensive behavior of a parent or teacher.

While the multiple channel approach is attractive for explaining a possible way for the acquisition of the semantics of certain concrete nouns we cannot use it when we want to model the language acquisition of older children, because they often acquire the meaning of new words by listening to explanations, i.e., by a single channel. We have modeled this type of language learning [Simha90,Fung91] by assuming that a child understands certain explanation patterns. These contain a handful of closed class words (such as determiners) which we assume are originally known, a number of words that are (could be) known through the multiple channel approach, and one completely unknown word per sentence. The patterns themselves are again assumed to be unambiguous.

The explanation pattern results in the creation of a memory structure (semantic network) such that the new word is grounded by the previously known words. It is not expected that the explanation patterns contain the full semantics of the new words, but that they connect every new word in some, possibly very indirect, way to words that are known perceptually, i.e., by the multiple channel approach. In this way we are trying to avoid the pitfalls that have been well described in the symbol grounding literature [Harnad90] that one cannot explain meaningless symbols with more and more meaningless symbols.

We have described [Simha90,Fung91] seven basic methods of "grounding" a symbol. The two most important ones are top-down and bottom-up grounding. As an example, imagine that a child is informed that

```
Cats and dogs are animals.
```

Assume in addition that the child has seen cats and dogs before, and the only unknown word in the above sentence is "animals". The above sentence creates a bottom-up connection from the words that are themselves perceptually grounded to the word "animals". Therefore, the word "animals" has now received perceptual reality in an indirect and **non-definitional** way.

3 Attribute Applicability

We will now discuss an extension of the single channel approach towards dealing with attributes. As noted in the introduction, we limit ourselves to attribute applicability. In this abstract we will use the term *attribute class* to describe what Lyons [Lyons77] calls

> ... a superordinate marker taken from the set M = { SEX, COLOR, AGE, SPECIES }... (p. 325)

and the terms *attribute, attribute value*, or just *value* to describe Lyon's

> ... subordinate marker, μ, specifying which particular location within the domain denoted by the superordinate marker is denoted by the subordinate marker." (ibid.).

We will discuss the specific problem of how to acquire the applicability of (1) attribute classes and (2) of attribute values to classes described by nouns. In other words, (1) how does a child learn that houses do not have taste, but they do have size? (2) How does a child learn that a house may be large or small, but it may not be sour, salty, or ferocious? This problem is more complicated than the previously discussed semantic acquisition problems, because the requisite information is normally not available in a single sentence.

In other words, the class of acquisition problems that we are dealing with now is more difficult than both of the classes that were previously mentioned, because, a child is **not** usually told that

```
A house can be red, and red is a color.
```

The simplistic single channel model used previously has to be replaced by a model that has some features of a discourse model in that it maintains some lexical and propositional information between the presentation of two sentences before it can draw the required conclusions. On the other hand, this is not a true discourse model, in that the child does not have to produce any utterances.

We have found that a good source of information about attribute applicability is contained in questions-answer pairs as they may occur when a parent or teacher poses a question and then answers it for the child. For instance, a parent might say

```
What is the color of the dog?
```

and then answer for the child by saying

```
The dog is brown.
```

This question-answer pair permits the derivation of three facts. The first fact comes out of the question alone, together with the pragmatic assumption that the teacher is not making category mistakes. The second and third facts come out of the answer:

1. The "attribute class" *color* is applicable to the class *dog.*

2. Brown is a color.

3. The attribute value *brown* is applicable to dogs.

One might want to dispute the necessity of the third fact. If we know that the attribute class A_c applies to the class C, and that A_v is a value of the attribute class A_c, then we should be able to conclude that A_v applies to C. Unfortunately this is not the case. In fact we are facing two difficulties.

1. An attribute value may not normally occur for a certain class, even if the attribute class in general does occur with this class. As an example, consider that dogs certainly have colors. Green is a color. Nevertheless, we cannot conclude that dogs are normally (or even sometimes) green. Note that we are not concerned about the fact that a dog could in principle be green if somebody chooses to throw a bucket of green paint at him. We are concerned about the fact that dogs in reality tend not to be green, and that children seem to acquire this type of knowledge. Therefore, we may not conclude:

$$\frac{\begin{array}{c}\text{applies-to}(A_c, C)\\ \text{value-of}(A_v, A_c)\end{array}}{\text{applies-to}(A_v, C)}$$

   unless we are interested in "being applicable in principle", independent of the reality of the world.

2. An attribute value may be applicable to a certain class, although one cannot say that its primary attribute class is applicable to that class. This might be due to metaphorical use of language. For example, the primary attribute class for the attribute value "sweet" is certainly "TASTE". In a metaphoric way we may apply the term "sweet" to a house, although the attribute class of taste is not applicable to houses. Therefore, we **cannot** conclude:

$$\frac{\begin{array}{c}\text{applies-to}(A_v, C)\\ \text{value-of}(A_v, A_c)\end{array}}{\text{applies-to}(A_c, C)}$$

   unless the attribute value A_v applies to a single class and is never used metaphorically.

In summary, one can say that the applicability of an attribute class to a class is not sufficient for guaranteeing that all attribute values are applicable to that class. The inapplicability of an attribute class to a class is not sufficient for guaranteeing that no attribute value of that attribute class applies to a class. Therefore, all three conclusions derived from the above question-answer pair are useful for the acquisition of the semantics of attributes.

4 Implementation

We have implemented the attribute applicability reasoning described in the previous section in the same environment as our previously described single and multiple channel language learning system

[Sen90,Simha90,Fung91]. However, this implementation is not connected to those two modules. An ATN parser processes incoming sentences that conform to a small set of sentence frames.

A combination of a semantic network, implemented with the SNePS system [Shapiro89], and a set of simple LISP constructs maintains information gained from processing these sentences. When a sequence of sentences is terminated by a pair consisting of a question and an answer to this question, then the system evaluates all the information contained in previously processed sentences. The system looks for sentences of the form

```
                <det> <noun> is <adj>
e.g.
                The house is red
```

and maintains a list of pairs (<noun> <adj>). In this way it learns that "houses can be red", etc. At any time the sequence of such sentences may be terminated by a pair consisting of a question and a presumed answer to that question. For example, the previous one sentence sequence might be terminated by the following question-answer pair:

```
What is the color of the dress?
It is red.
```

This pair furnishes the information that red is a color, and the system can now derive the fact that houses can have colors. The system will in fact report:

```
I found out.
House can have color.
```

5 Conclusions

We have presented our previous work in the area of Machine Learning of Natural Language. Two approaches, the single channel approach and the multiple channel approach were introduced. Then our current work on extending the single channel approach from single sentence processing to sentence sequence processing was discussed. It was shown that it is possible to derive the applicability of attribute classes and of attribute values to classes by processing question answer pairs. It was also shown that it is usually not sufficient to know the applicability of an attribute class to a class for concluding that all attribute values of that attribute class apply in the real world to that class. Finally, a prototype implementation of a system that derives attribute applicability information was presented.

ACKNOWLEDGEMENT

The described attribute applicability reasoner was implemented by Jeffrey Stavash. Examples are adapted from his report. Soon Ae Chun and Dan Devitt have commented helpfully on earlier drafts of this paper.

References

[Feldman90] Feldman, Jerome A., George Lakoff, Andreas Stolcke and Susan H. Weber. Miniature Language Acquisition: A touchstone for cognitive science. In *Twelfth Annual Conference of the Cognitive Science Society*, pages 686-693. Lawrence Erlbaum Associates, Publishers, Hillside, NJ, 1990.

[Fung91] Fung-A-Fat, Nigel. Acquisition of Language Semantics and Base Level Categories. Unpublished Master's Thesis, CIS Dept., NJIT, 1991.

[Harnad90] Harnad, Steven. The Symbol Grounding Problem. *Physica D 42*, pages 335-346, 1990.

[Lyons77] Lyons, John. Semantics. Cambridge University Press. New York, 1977.

[Markman90] Markman, Ellen M. Constraints Children Place on Word Meanings. *Cognitive Science*, 14(1), pp. 57-78, 1990.

[Powers89] Powers, David M. W., and C. C. R. Turk. Machine Learning of Natural Lan; Springer-Verlag, London/Berlin, 1989.

[Sen90] Sen, P. . Symbol Grounding. Unpublished Master's Thesis, CIS Dept., NJIT, 1990.

[Shapiro89] Shapiro, Stuart C., and The SNePS Implementation Group. SNePS-2.1 User's Manual. Department Of Computer Science, State University of New York at Buffalo. Buffalo, New York, 14260, 1989.

[Simha90] Simha, A. R. Symbol Grounding Paradigms–A Language Learning Perspective. Unpublished Master's Thesis, CIS Dept., NJIT, 1990.