

PAUL DE PALMA

Gonzaga University, Spokane, USA

E. JUDITH WEINER¹

Temple University, Philadelphia, USA

1. INTRODUCTION

In another study (Weiner and De Palma 1993), we determined that accessibility hierarchies among the meanings of polysemous words play an important role in the generation and comprehension of a category of simple riddles. In that paper and this one we restrict our attention to riddles that fit this definition:

A riddle consists of a single sentence question (RQ) followed by a single sentence answer (RA). The RQ presents a collection of attributes which can apply to more than one NP, thus allowing more than one answer to the question: a riddle answer (RA) and a straight answer. The RQ has been constructed in such a way as to lead the riddlee to prefer the straight answer over the RA.

Lexical ambiguity is at the center of riddles of this type. Consider this riddle:

- (1) RQ: What has a mouth and cannot eat?
RA: A river.

Here the riddler intends by *mouth* the inanimate meaning but the sentence is constructed so that the animate one is the more natural reading. In our 1993 paper, we showed how the existence of accessibility hierarchies could account for this preference.

We now turn our attention to the need to build this mechanism into any knowledge representation language that hopes to capture the full subtlety of natural language since it is our contention that riddles violate the rules of

normal discourse and thus represent a useful way to approach the study of these rules. To that end, we present a knowledge representation plan along with an algorithm that uses the representation in order to generate riddles. Although the representational structures that we use are in the style of KL ONE (Brachman and Schmolze), this is purely a convenience. Accessibility hierarchies must be built into any system which can process natural language as well as one which can play the riddling game.

2. ACCESSIBILITY HIERARCHIES

Cognitive psychologists have long recognized that people form taxonomic categories (Rosch 1978) with some members being more typical instances of those categories than others. This graded structure is not limited to taxonomic categories but seems to include such unlikely possibilities as formal categories like the category of odd numbers and that of squares (Armstrong, Gleitman, and Gleitman 1983) and linguistic categories for phones, phonemes and syntactic structures (Lakoff 1986). In recent years, researchers have shown that categories are not structurally invariant, but are, in fact, demonstrably unstable (Barsalou 1987). Their graded structure varies with such factors as linguistic context and point of view and even differs for the same individual over time.

The formation of ad hoc categories to accomplish specific goals (Barsalou 1983) is another area of instability in human category formation. For example, the category "things-to-take-out-of-a-burning-house" might include subordinate categories like

¹The order of the names of the two authors is arbitrary. This paper represents the joint work of both authors.

"children," "jewels," "paintings," and "portable tv's" (Murphy and Medin 1985) and is formed only when one's house is burning or during a discussion like this one. Ad hoc categories, once formed, function similarly to more traditional categories. As we show later in this paper, ad hoc category formation is an important component in the generation or solution of riddles.

A model that is to account for the human tendency to form categories must account for both the stable and the unstable aspects. Barsalou's approach to the instability in categories is to recognize the existence of both context-independent and context-dependent information in long-term memory, where it is arranged as interrelated, continuous knowledge. It is the context-independent information that is most likely to be shared by a large number of individuals within a speech community. Its activation is obligatory. When one thinks of robins, for example, "red-breasted" springs to mind whereas "poisonous" is triggered by rattlesnakes. Context dependent information, by contrast, is accessed only within a relevant context. So "hard" may be triggered by ice while discussing fall. The instability of categories is accounted for by different information for a given category being compiled in working memory at different times depending on the situation. Some information, e.g., context-independent information, is more accessible than other information.

We have extended this model (Weiner and De Palma 1993) to explain the tendency of people to think of the mouth of a person before mouth of a river in (1) above. Given the presumed universality of certain principles governing categorization, it seems likely that, in context-neutral situations such as (1), ambiguous words form ad hoc category-like structures of their multiple meanings onto which an accessibility hierarchy is imposed. For example, in (1), there is a category-like structure corresponding to the phonemic realization of the word mouth to which the different meanings belong; in (1), one thinks of

the mouth of a person before the mouth of a river.

3. THE KNOWLEDGE BASE

We thus offer our exposition of the structure that underlies the kind of lexical ambiguity found in riddles as linguistic evidence for the epistemological requirements of a knowledge representation system which can support both normal discourse and riddles. Riddles will use the knowledge in one way; normal discourse will use it in another. The representation will remain the same; only the algorithms will differ.

Consider Figure 1, a knowledge-base fragment in the style of KL-ONE that contains the information necessary to generate or solve riddle (1). The KL-ONE entities most relevant to this discussion are Concepts (diagrammatically represented by ellipses) and RoleSets (represented by encircled squares). The Concept is the primary representational entity. For us, it represents the category of objects indicated by the Concept name. Thus, in Figure 1, Concepts stand for the category RIVER_MOUTH, the category ANIMATE_MOUTH, and so on. Concepts are connected to one another by superC links, represented in the figures by double arrows. A superC link indicates that the subordinate Concept (subConcept) stands in an inheritance and subsumption relationship with the superordinate Concept (superConcept). (The higher Concept subsumes the lower one; the lower one inherits from the higher one). Thus, PERSON_MOUTH is an ANIMATE_MOUTH and a MOUTH.

In our knowledge base, RoleSets represent predicates of a Concept, the fillers of which, known as Value Restrictions (v/r's), are themselves Concepts. So PERSON_MOUTH has a RoleSet "function" with the filler EAT, meaning in our representation that a function of a person's mouth is to eat. (Of course there are others not shown here).

Further, each RoleSet filler has a number restriction represented by two numbers within parentheses. These represent the lower and upper bounds on the number of fillers for a




Figure 1

KL-ONE-like representation of a portion
of the knowledge needed to generate or solve:
What has a mouth and does not speak?

given RoleSet. In Figure 1, we have arbitrarily estimated that people's mouths have a minimum of 5 and a maximum of 5 functions.

Notice that every Concept has a diamond-shaped symbol associated with it. This symbol is not part of the KL-ONE language. We are introducing it here as a new primitive, Lexical, which contains lexical information about a Concept. For our purposes, Lexical contains the phonemic representation of a Concept (although, for simplicity in this figure, only certain phonemic representations are actually provided). This arrangement allows us to acknowledge the relationship between a Concept and the word used to name the Concept without asserting that they are the same thing, separating meanings of polysemous words from their phonemic representation.

As discussed above, ambiguous (polysemous, homophonous) words can form ad hoc category-like structures of their multiple meanings. Thus, we can have a superConcept MOUTH, a category of polysemous words, with subConcepts ANIMATE-MOUTH and INANIMATE-MOUTH. We recognize the probability that in the case of ambiguous forms with a choice of animate vs. inanimate meaning, the animate one is thought of before the inanimate one (Weiner and De Palma 1993). So the ideas encoded in Figure 1, although not explicitly spelled out with respect to accessibility, are based on the assumption that, in context-independent situations, people tend to think of animate things before they think of inanimate ones.

In riddle (2),

- (2) RQ: What has four legs and only one foot?
RA: A bed.

we model the riddling process by assuming that the phrase four legs causes the formation of an ad hoc category "four legged thing." A representation of a portion of the knowledge needed to generate or solve riddle (2) will be given in a future paper.

4. THE ALGORITHM

The following algorithm refers to Figure 1 and will generate riddle (1). The algorithm requires three functions:

1. FindHoms(HC1,HC2,C1,C2) - searches the knowledge base for two homophonous Concepts, HC1 and HC2 where HC1 and HC2 are the value restrictions of two Concepts' RoleSets. Call these Concepts C1 and C2. C1 must contain the more accessible (i.e., in these examples, context-independent, animate) concept. For example, after an application of FindHoms(HC1,HC2,C1,C2), on the KB fragment contained in Figure 1, the variables would look like this:

```
HC1 <--- PERSON_MOUTH
C1 <--- PERSON
HC2 <--- RIVER_MOUTH
C2 <--- RIVER
```

Note that HC1 contains PERSON_MOUTH, a value restriction of C1 (PERSON), HC2 contains RIVER_MOUTH, a value restriction of C2 (RIVER) and the Concept in C1 (PERSON) is a more accessible Concept than the one in C2 (RIVER).

2. Lex(A,B) - returns in B the word by which a Concept, A, is known. Remember that the phonemic representation of this word is contained in "Lexical" (represented in the figure by the diamond shape) for each concept. For example, Lex(RIVER_MOUTH,B) returns /mawθ/ in B.

3. MisMatch(C1,C2,HC1,HC2,Type,RSVR) - examines the knowledge base (KB) for a mismatch of the following type:

HC1 has a RoleSet value restriction (RSVR) that HC2 does have. In Figure 1, this RSVR for HC1 would be EAT. Mismatch returns this in RSVR. Thus, using Figure 1, Mismatch would return EAT in RSVR. Note that HC1 is more accessible than HC2 by virtue of being animate.

The algorithm, then, looks like this:

```

Riddle_Gen()
  FindHoms (HC1,HC2,C1,C2);
  MisMatch(C1,C2,HC1,HC2,Type,RSVR);
  Print "What has Lex(HC1) and ~Lex(RSVR)?"
End.

```

It should be noted that, in the interest of simplicity, we have conflated the issues involved in generating or solving riddles. Once you know the heuristic with which riddles of the type considered in this paper are constructed and have created a KB of Concepts, generation is a simple matter. Solution, of course, is the inverse of this algorithm.

5. CONCLUSIONS

Our examples in this paper use KL-ONE as a convenient model of a knowledge representation system. We propose the addition of accessibility as an important epistemological primitive to the KL-ONE system since it appears critical to build this factor into any knowledge base which can both support a system for natural language processing and be used for certain kinds of humor. Our work also highlights other requirements for knowledge representation systems capable of supporting natural language:

1. Links between the phonemic representation of linguistic entities and their associated concepts (lexical)
2. The necessity of representing homophonous categories
3. The ability to form ad hoc categories such as those based on homophonous phrases

REFERENCES

1. Armstrong, S., L. Gleitman, and H. Gleitman 1983 What some concepts might not be. *Cognition* 13: 263-308.
2. Barsalou, L.W. 1983 Ad hoc categories, *Memory and Cognition* 11(3): 211-227.
3. Barsalou, L.W. 1987 The instability of graded structure: implications for the nature of

concepts. In U. Neisser (ed.), *Concepts and conceptual development: ecological and intellectual factors in categorization*. New York: Cambridge University Press, 101-140.

4. Brachman, R.J. and J.G. Schmolze 1985 An overview of the KL-ONE knowledge representation system, *Cognitive Science* 9: 171-216.
5. Lakoff, G. 1986 *Women, fire and dangerous things: What categories tell us about the nature of thought*. Chicago: University of Chicago Press.
6. Murphy, G. and D. Medin 1985 The role of theories in conceptual coherence, *Psychological Review* 92(3): 289-316.
7. Peters, S., S. Shapiro and W. Rapaport 1988 Flexible natural language processing and Roschian category theory, *Proceedings of the 10th Annual Conference of the Cognitive Science Society*: 125-131.
8. Rosch, E.H. 1978 Principles of categorization. In E. Rosch and B.B. Lloyd (eds.), *Cognition and Categorization*. Hillsdale, NJ: Lawrence Erlbaum Associates, 27-48.
9. Weiner, E. J. and P. De Palma 1993. Some pragmatic features of lexical ambiguity and simple riddles. *Language and Communication* (in press).

ACKNOWLEDGMENTS

We would like to thank David Weiner for his invaluable advice and assistance during the preparation of the final manuscript.