

A Knowledge Representation Approach to Understanding Metaphors

E. Judith Weiner

Computer and Information Sciences
Computer Activity Building
Temple University
Philadelphia, PA 19122

This study represents an exploration of the phenomenon of non-literal language (“metaphors”) and an approach that lends itself to computational modeling. Ortony’s theories of the way in which salience and asymmetry function in human metaphor processing are explored and expanded on the basis of numerous examples. A number of factors appear to be interacting in the metaphor comprehension process. In addition to salience and asymmetry, of major importance are incongruity, hyperbolicity, inexpressibility, prototypicality, and probable value range. Central to the model is a knowledge representation system incorporating these factors and allowing for the manner in which they interact. A version of KL-ONE (with small revisions) is used for this purpose.

1. Introduction

One can hardly fail to notice the flurry of intellectual activity that currently surrounds the understanding of the use of figurative language. The interest is multidisciplinary – linguistics, psychology, philosophy, education, to name a few of the more active disciplines. The reason, which anyone writing on the subject hastens to point out, is that the observation of natural speech demonstrates clearly that it is rarely confined to the strictly literal. Figurative language is not merely an ornament of the poet but abounds in the everyday speech of everyday people and as such is a legitimate area of inquiry for researchers – in any discipline – who are concerned with understanding natural language. The interest in metaphor in computer understanding of natural languages stems from this same source. It is well understood that people, when conversing with machines, can no more be constrained to literal language than they can be expected to be long contented, within the confines of a synthetic language.

2. Scope of the Study

The heading “figurative language” comprises the traditional figures of speech known as synecdoche, metonymy, hyperbole, personification, irony, etc., as well as the more common metaphor and simile. I am going to focus here on these latter two in order to narrow my

view in the hope of achieving some depth and also because of a belief that the other figures may operate under similar principles. Except where noted, I will use the term “metaphor” in referring to both similes and metaphors. Please note that this does not imply that I am taking the position that metaphors and similes are the same; in fact, there is some evidence that they function differently from one another. At the least, it seems possible that the distinction between these two is more than the traditional one of implicitness versus explicitness since there are instances of metaphors that sound strange when “transformed” into similes and vice versa. I therefore am using the term “metaphor” in a very loose way to cover the area metaphors and similes have in common (for example, the similarity in the figurative reads of *John is an animal* and *John is like an animal*), without pausing at this time to delve into its exact nature and ignoring for the moment the apparent differences.

To start, I will work only with isolated sentences of the form

(1) A is (like) B.

In sentences of this form, A is commonly referred to as the “topic”, the B term as the “vehicle”. That which they have in common is called the “ground”. In a sentence like

Copyright 1984 by the Association for Computational Linguistics. Permission to copy without fee all or part of this material is granted provided that the copies are not made for direct commercial advantage and the CL reference and this copyright notice are included on the first page. To copy otherwise, or to republish, requires a fee and/or specific permission.

0362-613X/84/010001-14\$03.00

(2) Billboards are like warts

then, the topic is *billboards*, the vehicle *warts* and the ground *ugliness* and (perhaps) *prominence*.

In restricting this study to sentences of the form (1), my motives here again are to constrain the unwieldiness of the subject. Of course this rules out a large body of possible metaphorical utterances of other forms. Many of these, however, if confined to one sentence, could be restated in the form of (1) with no significant loss of meaning. I will not discuss them here. I will, however, have something to say later about the larger linguistic context (discourse.)

The typed word, the presumed form of input of natural language to a computer until such time as actual speech understanding systems develop sufficiently, imposes limitations of its own on the scope of any language processing system. The most obvious is, of course, that variation in intonation of the input is limited to its most "neutral" pattern; prosodic features must largely be ignored. (A certain amount of emphasis or contrastive stress may be obtained by underlining, but the study of this should be considered separately.) Another, perhaps more relevant consideration, is the use of a space to separate parts of what must be considered a single lexical item, e.g., *blind alley*. Historically, this was undoubtedly a metaphor (and a candidate for this study); today it is most probably interpreted as a single unit. Although most native speakers of English would classify it as an idiom through an awareness of the written form and the fact that even in the spoken version the component parts are clearly recognized, there should be no attempt to componentially process such forms.

3. Salience

One of the most useful notions for modeling metaphoric understanding is that of salience (Ortony 1979a), which Ortony takes to mean an estimation of "prominence of a particular attribute with respect to a concept to which it does or could apply." (p. 162.) He later speaks of "predicates" rather than attributes" (1979b, p. 191): "A predicate can be attributed to, or predicated of, something. It can represent knowledge, a belief, or an attitude about or toward something." I too prefer the flexibility of "predicate" and shall follow Ortony in the use of this term.

The notion of salience makes use of the apparent fact that metaphorical statements are asymmetric:

(2) Billboards are like warts

means something different from

(3) Warts are like billboards

Ortony's explanation is that in isolated sentences of the form

(1) A is (like) B

those predicates that have high salience in B and low salience in A are the ones being considered in the metaphor.¹ The effect is one of raising the salience of these predicates in A. Thus, sentence (2) is generally understood as meaning that billboards are ugly, whereas in (3), those predicates that have high salience for billboards (but low salience in warts) – for example, prominence – are attributed to warts (that is, the salience is raised). An additional requirement is that there be high salient predicates of B that cannot apply to A.²

My working definition of salience includes the assumption of a taxonomic structure of concepts with the most general at the top and the most specific at the bottom. Figure 1 provides an illustration of a simple taxonomy.

I define a salient predicate of a concept as one that implies inherent prominence (for example, saturation of color, largeness of size, etc.) or else is definitional in that it entails a concept's separation from others in the hierarchy (for example, the dog's domesticity separates it from the wolf). My notion of salience parallels that of Tversky (1977) in that intensity and diagnosticity are the critical factors, but it makes the addition of hierarchical organization to facilitate diagnosticity.³ Context, both linguistic and extralinguistic, is of course a major contributing factor too, but it is outside the scope of this study. The way in which these factors interrelate is a fertile area for psychological research.

4. Prototypicality

Another valuable contribution coming out of cognitive psychology is prototype theory (Rosch 1973, Rosch and Mervis 1975), which holds that a concept may belong to a category even if it is somewhat atypical in terms of the predicates usually (typically) associated with members of that category. A chicken is a bird even though it can't really fly. Here, *bird* refers to some prototype from which chicken represents a departure.

In terms of metaphors, there is much value in including prototype theory in any model. For example,

(4) Mary's cheeks are like apples

¹ Ortony acknowledges that the attributes may be similar, not identical, in the vehicle and topic (1979a, p. 167). While I am in agreement with Ortony, I will, for the purposes of this paper, make the assumption of predicate identity.

² In testing this hypothesis experimentally, Gentner (1980) showed that salience did not appear to be a relevant mechanism in metaphor processing. This seems to me, however, to be partially a result of how salience was measured and the need for a clearer analysis of how metaphoric interpretation proceeds. Salience, properly defined, may provide a necessary but not sufficient explanation.

³ Conklin and McDonald (1982) have used salience as a means of solving the selection problem in natural language generation using KL-ONE as the representation language.

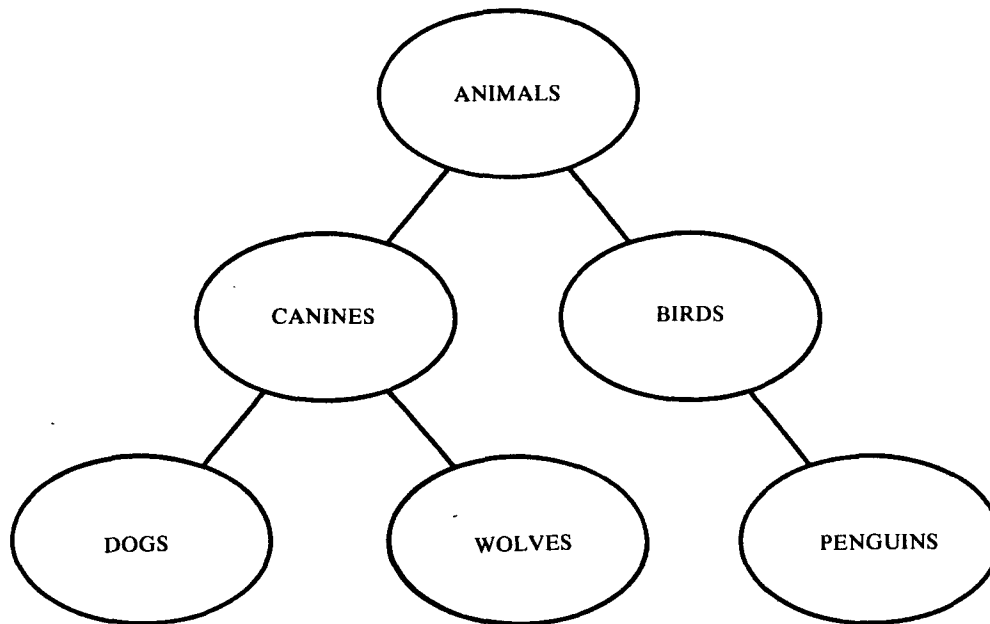


Figure 1.

would probably mean to most people that Mary's cheeks are round and red. A different interpretation would be obtained if the concept of round, red apple were replaced by a withered, rotten one or even, for that matter, by a green one. For communication to take place among people in a speech community, some sort of prototypicality considerations are essential.

There is another way in which prototypicality might figure in a discussion of metaphors. It seems probable that the B term in

(1) A is like B

represents the *epitome* of the predicate(s) that are true of A and of interest in a given utterance, that is, B is the prototypical representative of these predicates. (Tversky (1977) observed that the B term is the more prototypical of the two in literal sentences.) If the sentence reverses A and B, then A becomes the prototype of (probably) different predicates. The vehicle of choice should be one in which the cluster of predicates is (ideally) uniquely appropriate, prototypical, and therefore also salient. For example,

(5) A hose is like a snake

Snakes are typically, even classically, the ultimate in long, narrow, coiledness; these characteristics can be thought of as distinguishing snakes from other members of the category ANIMALS. Sentence (5) draws the reader's attention to these (perhaps) slightly less salient qualities of a hose.

I have intentionally limited myself to sentences taken out of the discourse context. One of the benefits of doing this is that there is considerable context

within the sentence itself that can influence its interpretation. Consider the following pair of sentences

(6) My cat's tail is like a carrot

(7) John's hair is like a carrot

Without adding any context, it is unlikely that a person would miss the fact that the relevant salient predicate of (6) is shape (and perhaps color) and that of (7) is color. Our knowledge of prototypical cats with prototypical cat tail shapes and colors and prototypical hair shapes and colors leads us to the right conclusion. If no additional information is available, then it wouldn't be likely that John was wearing a pony tail. On the other hand

(8) John's nose is like a carrot

again, taken out of context, would indicate a comment about shape. From these examples it should be clear that prototypicality considerations are relevant to both topic and vehicle.

5. Prototypicality and Possibility

How then, does this relate to prototype theory? The relationship appears to be that, in the *prototypical* tail, nose, or hair, certain predicates are probable. These are the ones most likely to match those salient in B. Thus in order to process metaphors, it is necessary to know, in addition to the nature of the prototype, a range of probable values for a given predicate. This would facilitate Ortony's determination (1979a, p. 173) of "whether any gross incompatibility would result by applying the predicate in question to the

concept." This range can help determine whether the statement is literally true or not. For example,

(9) John's hands are like ice.

If a range of possible temperatures were built into the representation for human hands, it would be known that John's hands could not possibly be literally as cold as ice (there could not be an actual equivalence of temperature in hands and in ice).⁴ The figurative interpretation would thus prevail.

Prototypicality and range of possible values appear to operate throughout language as a whole; they are not confined to figurative usage. Labov (1973) reported on an experiment on the denotation of the word *cup*. He gave subjects pictures of cuplike objects to identify and observed differing percentages of the use of *cup* as form and function were varied. He concluded the existence of an invariant core (corresponding to my use of prototype) as well as a range of deviations through which recognition still occurred, albeit at lower percentages. The inclusion of these elements in the knowledge representation of a system for understanding natural language is therefore broadly motivated.

6. Metaphors as Hyperboles

All metaphors are hyperbolic in a sense. They seem to say: the predicates A shares with B are in A so extreme that they can only be expressed by relating them to some object in which they are epitomized, that is, B. In

(9) John's hands are like ice,

the exaggeration is apparent. This is evident because of the range of possible temperature values known to be actually attributable to human hands. The sentence

(7) John's hair is like a carrot.

is less metaphoric (consequently more literal) in this sense. It is *possible* that hair could be the same color as the prototypical carrot, but the probability is low. Consequently, the following sequence does not seem absurd:

(10) John's hair is like a carrot.
Is it *really* that color?

whereas

(11) John's hands are like ice.
Are they *really* that cold?

would lead one to think the response peculiar at the least. On some scale of metaphoricity then, (10) is less metaphoric than (11). Since a hose may in fact be as long, narrow, coiled as a snake,

(5) A hose is like a snake

is the least metaphoric of the examples given. This supports Ortony's claim that high salient predicates of A matched with high salient predicates of B make for a literal statement. A response of *Is it really that long, narrow, coiled?* could easily be followed by an unqualified response of "Yes."

Ortony (1975) denies the possibility of the ground consisting merely of a single predicate. "People simply do not use metaphors to transfer *one* characteristic, even if it is a distinctive one, when there is a ready literal way of making the point." (p. 50.) Sentence (9), however, provides a fairly good counterexample to this claim. Here, the hyperbolic nature of the metaphor rather than the size of the ground provides the incentive for its use.

7. Taxonomic Structure and Incongruity

The conclusion which should be drawn from this discussion is that all of the above factors must be brought to bear in an analysis of metaphor understanding. The result of using these measures will be an isolation of those predicates of B that are true of A and the establishment of a relative degree of metaphoricity within the sentence context.

Given a corpus of sentences of the form

(1) A is (like) B,

some will appear to be literal similarity statements; others will appear to be metaphors.

(12) John is like his father.

(13) John is like a snake.

(14) John is like a black box.

In (12), the sentence appears to be a literal comparison. John shares certain characteristics with his father. John and his father are already known to be similar on the basis of the fact that they are members of the same superordinate category (males) or are already known to participate in a relationship to one another (father-son). There is no element of surprise or incongruity in statement (12). As Ortony has proposed, high salient predicates of B are also high salient predicates of A. This is Ortony's criterion for a literal similarity statement as opposed to a metaphor. I think, however, that viewing this phenomenon from the perspective of category membership, relationship and consequent incongruity will shed more light on its computation representation.

There do seem to be metaphorical statements in which there is matching of high salient predicates in *both* the vehicle and the topic. If sentence

(8) John's nose is like a carrot

were uttered by one of John's friends to another, it would not represent new information. It would proba-

⁴ If John is assumed to be living and context does not indicate the possibility of frostbite conditions.

bly produce laughter because of the hierarchical incongruity but beyond this should have no more of an effect than a literal paraphrase of the same sentence. In terms of those predicates that can be shared, those peculiar to the vehicle's hierarchical position – for example, in (8), a carrot's being a root crop – seem the least likely candidates.

The position in a taxonomic hierarchy is important in another way. Consider

(15) Penguins are like wolves.

(16) Dogs are like wolves.

(16) is a similarity statement; (15), assuming an interpretation can be found, is a metaphor. Figure 1 provides a possible explanation: In terms of this diagram, dogs and wolves are siblings but penguins and wolves are not. The latter relationship is more distant, the sentence more metaphorical. Although distance metrics are notoriously difficult to pin down, fairly clear cases like this one indicate that members of categories at a level of abstraction close to that of Rosch's basic level categories (1973) and with a shared superordinate can be considered in some way closer than those without a shared one. Empirical research on other categorical relationships may provide additional support for this approach. Although less is known about the relationships in the examples that follow, since they no longer deal with a basic level category and its immediate superordinate, future research may shed some light on this phenomenon as well.

For example, that incongruity may be a factor of "goodness" of metaphor seems to be illustrated by (13) and (14). Metaphors may be judged by the apparent unlikelihood of an A B juxtaposition. This of

course is under the provision that the metaphor is understandable to the hearer. Thus (14), if understood to mean that John is somehow unknowable, is better than (13) in which John is thought to be sneaky. The difference may be that in (13) A and B are of a shared superordinate category (ANIMALS), whereas in (14) they are categorically more remote.

Incongruity appears to be the reason that the best metaphors often produce a smile by the hearer as they are comprehended. ("How unlikely yet how apt" may be the attendant thought.) So there is a connection between metaphor and humor, or the intelligence and wit of the speaker who first utters a good, novel metaphor.

8. Number of Shared Predicates

In addition to incongruity, another attribute of metaphoric quality is the number of shared characteristics under incongruous circumstances. The more the better so long as incongruity is maintained:

(17) Jane's eyes are like stars

Although (17) is somewhat hackneyed, its survival and wide use may be a result of the fact that it is good in the sense just described, that is, incongruity plus multiplicity of shared predicates (twinkliness, brightness, beauty, clarity, etc.).

Some writers have talked of the magic of metaphors, the idea that the whole is equal to more than the sum of its parts (Verbrugge 1980). While there may be other factors involved than those I have mentioned above, I believe ultimately that metaphors can at least theoretically be accounted for formally and without appeal to the supernatural.

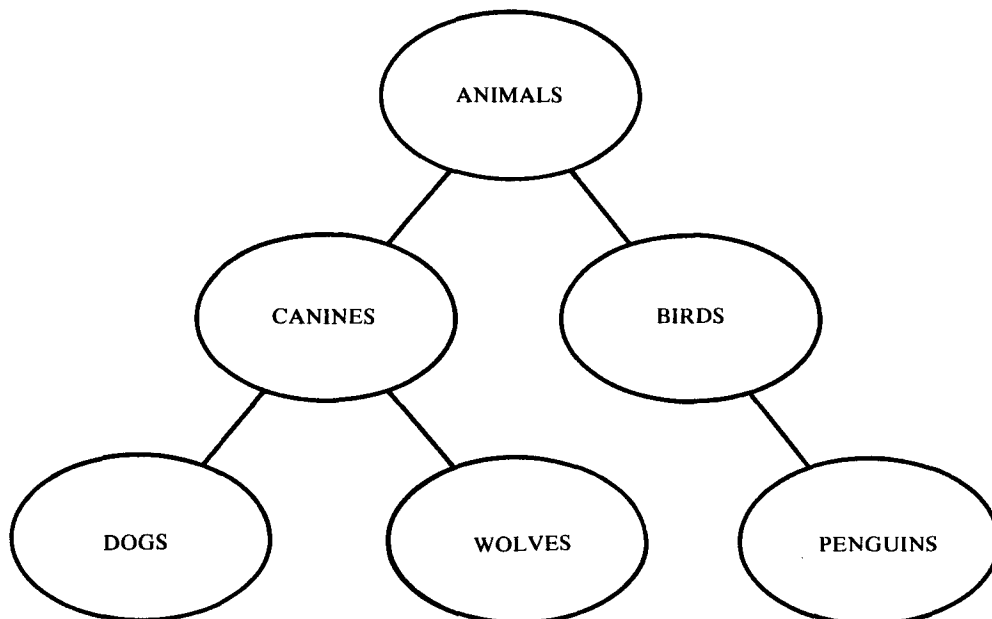


Figure 1.

9. Expressing the Inexpressible

Why do people use metaphors? Inexpressibility in literal terms seems to provide part of the answer.⁵ Hyperbole as discussed above is one way that this is overcome. When an extreme in terms of predicates is unavailable to the speaker, he or she may search other domains (categories) for something that epitomizes the desired quality. Metaphor makes the difficult to express possible to express.

This provides an explanation of the fact that human emotions (love, hate, etc.) are often described metaphorically, that is, the less well-understood in terms of the more thoroughly understood. It also helps explain the observation that more abstract concepts with few if any high salient predicates are used for A while more concrete concepts (with predicates of high salience) are used for B. There are many examples of metaphors for love, friendship, etc. in which B is a concrete concept having usable salient predicates.

Metaphors appear commonly with regard to people to express something about their more abstract personal characteristics (personality, character, value) rather than the more easily stated physical attributes.⁶ Thus, when searching for the meaning of such metaphors, all other things being equal, one should generally exhaust those first. When one says,

(18) You are my sunshine,

one is not attributing yellowness to the addressee but rather those characteristics that are both more abstract and possible to attribute to a human – that is, warmth, brightness, cheerfulness, etc. If such characteristics are unavailable or not salient in the B term, then other, physical ones are used in the interpretation.

(19) Jane is a string bean,

for example, makes a statement about Jane's shape.

It should be noted at this point that I am not making any claims about the historical primacy of physical over more abstract predicates (as do Lakoff and Johnson 1980). I am not assuming (nor am I contradicting) the possibility that this is true. The process by which it came to pass that we can now say "John is a bright person" or "That is a bright light" is not the issue. It does seem reasonable that the abstract was at one time a metaphor based on the physical, but I don't know what the historical evidence is for this. My analysis of metaphor is strictly a synchronic one. I am

⁵ Ortony's inexpressibility thesis (1975) deals with transferring from vehicle to topic "characteristics which are unnameable."

⁶ Carbonell (1981) has proposed an invariance hierarchy for explaining this phenomenon in which physical descriptors occupy a relatively low position. An interpretation is obtained by searching downward through the hierarchy and stopping when knowledge common to A and B is encountered.

saying that this is how the language at any given historical stage can be perceived as operating.

10. A Word about Context

In extending the context of a metaphorical sentence to include the surrounding linguistic environment, Ortony (1979a) has suggested that the effect of the linguistic environment is to raise the salience of certain predicates. To analyze this notion, consider the following pair:

(20a) Look how the highway curves

(20b) It's just like a snake.

(20a) can be thought of as raising the salience of curviness in that sentence. Compare this pair to

(21) This highway is like a snake,

that is, the isolated sentence. The difference between (20) and (21) is that in (20) the discourse phenomenon of focus (Grosz 1977, 1981) is operating. Attention is focused (and salience consequently raised) to the snakelike curviness of the highway. Note however that the other attributes that have high salience in (20) can also be applied metaphorically. It is just that one particularly snakelike attribute is *highlighted* in the pair.

The following linguistic environment is also of importance and

(22) Look at how he eats.

Isn't John a pig!

(23) Isn't John a pig!

Look at how he eats.

Metaphors in literature, especially poetry, represent a possible limit to which a computational model might aspire since poets are experts in the novel use of language and in explicating human experiences. But it is true that these sources should not be overlooked just because they seem to present difficult problems; a language understander should at least theoretically be able to understand poetic metaphors. This poem by Emily Brontë⁷ is metaphorically fairly straightforward as poems go and a good illustration of how metaphors in context can function.

Love and Friendship

Love is like the wild rose-briar,
Friendship like the holly tree—
The holly is dark when the rose-briar blooms
But which will bloom most constantly?

The wild rose-briar is sweet in spring,
Its summer blossoms scent the air;
Yet wait till winter comes again

⁷ *The Mentor Book of Major British Poets*, edited by Oscar Williams, 1963.

And who will call the wild-briar fair?
 Then scorn the silly rose-wreath now
 And deck thee with the holly's sheen,
 That when December blights they brow
 He still may leave thy garland green.

The first two lines represent two similes of the sort that I have been discussing; in each case an abstract concept is juxtaposed with a concrete one. The remainder of the poem proceeds to describe the ground of the similes. Certainly the ground consists of some of the low salient predicates of wild rose-briars and holly bushes. This seems to expand Ortony's thesis that high salient predicates of B are the only ones to be considered. Those of low salience are also eligible. Including these predicates in the context of the poem serves to raise the salience of some predicates the reader may not have even had in his/her conceptual representation prior to reading the poem. Love and wild rose-briars do make an incongruous twosome as do friendship and holly trees. The success of the metaphor rests on this and on the size of the ground. The best metaphors are those presented by the best poets, those in which a vaguely understood experience is clarified through the predicates, salient or otherwise, of the B terms. Ortony's compactness thesis (1975), which allows metaphors to cause the transfer of features or characteristics as a "chunk" from vehicle to topic, does not account for this type of metaphoric discourse.

11. The Semantic Net Approach

From the above discussion it is obvious that some sort of conceptual representation underlies human metaphor understanding; at some level people know the predicates of concepts and presumably something about their organization in terms of a generality hierarchy. Our "knowledge" undoubtedly includes, in addition to what might be labeled general knowledge, the values and beliefs of our speech community. Although the computational representation of such a base is a formidable task, the purpose of this paper is to delineate its nature and boundaries. It is assumed that the implementation of a base, in actuality, is a separable task.

KL-ONE is one of a number of extant knowledge representation languages that allow Concepts to be arranged in a generality hierarchy with the characteristic that properties of the more general Concepts are inheritable by the more specific ones. It furthermore provides for the Concept to be represented as a structured object, allowing one, in effect, to get inside the Concept and to see its relationship to other Concepts. In this discussion, I will assume that the features of KL-ONE are available for use in metaphor processing.

The KL-ONE entities most relevant to this study are Concepts (diagrammatically represented by ellipses)

and Roles (represented by encircled squares).⁸ The structured Concept is the primary representational entity. A Role is internal to a Concept; it can be a part (for example, a hand is a Role of the Concept BODY) or what is commonly called an attribute (for example, a PERSON – the Concept – has habits – a Role). I will use Roles to represent predicates as they have been described above. The hierarchical classification aspect of KL-ONE allows lower, less general, Concepts to inherit structured description from those in an ancestral relationship to them.

At the topmost level are the most general Concepts, called Generic Concepts (GCs). As one proceeds downward through the network, one encounters more and more specific GCs. At the lowest level is knowledge about an individual, called an Individual Concept (IC) (see Figure 2). ANIMAL (a GC of a higher level) passes down to PERSON (a GC of a lower level) all of its Roles and the interrelationships among them. So if ANIMALS have noses then so do PEOPLE. PERSON in turn passes to JOHN (an IC, shaded in Figure 2) those Roles PERSON got from ANIMAL as well as those unique to PERSON.⁹ The Value and Satisfies links indicate the relationship between a Role of an IC (called an IRole) and its parent role.

Since the kind of knowledge necessary for metaphor processing must include the beliefs and cultural values of the members of the community for which the system is being designed, in the representation of a prototype, some stereotyping is inevitable. That snakes are frightening and perhaps evil creatures is a commonly held opinion, although of course this is not true of all snakes. Metaphors seem to tap these kinds of generalizations in their insistence on prototypes, and they do seem necessary for understanding to take place.

(24) The whip lay coiled on the ground like a snake.¹⁰

⁸ The local internal structure of every Concept is made up of Roles and RoleSet Relations (RSRs). This discussion will not describe the functioning of RSRs but will focus on Roles since they are adequate to support this theory at its present stage of development. Furthermore, RSRs are less clearly understood at the current state of KL-ONE's design than are Roles. (I use the term "Roles" instead of "RoleSets" for discursive simplicity.) Those readers who have an interest in a more detailed description of the knowledge representation language KL-ONE should consult Brachman (1978, 1979) and Schmolze and Brachman (1982).

⁹ A Concept description in KL-ONE represents the intension of the Concept and there is a clear distinction drawn between the intension and its extension (the Concept in the real world). There is also a distinction drawn between definitions and assertions. My examples will deal only with definitions. I will not discuss issues related to the definability of Concepts here. For a treatment of this, see Cohen (1982b).

¹⁰ My thanks to David Weiner (personal communication) for this example.

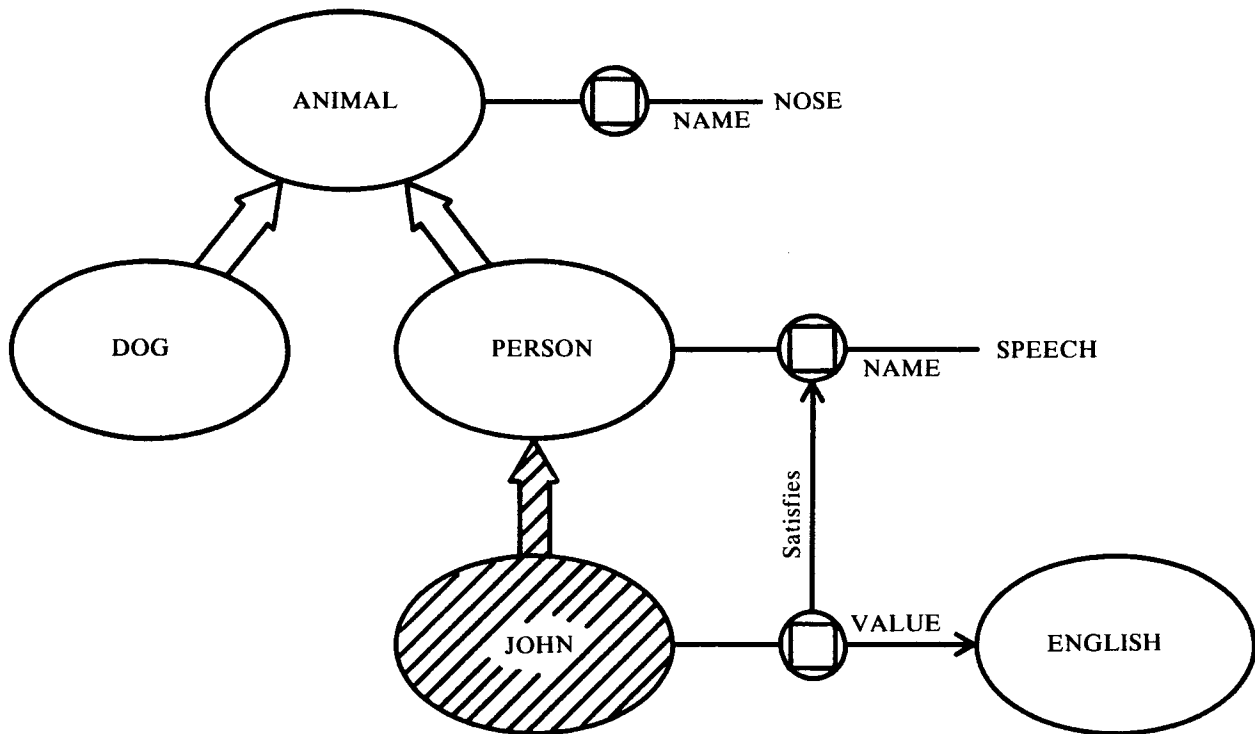


Figure 2.

That the more affective (here, negative) aspects of peoples' impressions of snakes are of importance, and not simply shape or the more physical characteristics, seems evident by making a comparison

- (25) The whip lay coiled on the ground like a strand of spaghetti.

The fact that a snake might strike and inflict harm from a coiled position (even though many types of snakes would not) shows how generalizations seem to operate. The menacing nature of the whip is highlighted in (24), but in (25) that aspect of whips is much less important. In fact, it can be considered to be "negated" by (25).

It should be noted at this point that one network alone may not suffice in processing metaphors or, for that matter, in processing other language phenomena. The existence of sublanguages is generally accepted within the field of linguistics: there are technical sublanguages for technical fields, for example. The representation must parallel and support this phenomenon. An individual's style of speech also changes according to the social setting; a person surely has more than one style. At the least, there is careful speech and casual speech. In lecturing to a class, one would use a more careful variety than in chatting with one's friends. I propose that there also exist sub-knowledge networks to support different styles of speech. Again, at the least there would be a careful and a casual variety. Formal situations would favor the careful; infor-

mal situations the casual. The other could then be a reasonable second choice.

As an example, a veterinarian would have one representation of the animal kingdom for use on the job and one for home use. The careful (or more technical) one might look like an expanded and detailed version of Figure 2. The sentence *John is an animal* might receive one interpretation (the literal one) at work where, due to the nature of the representation (and of course the context), the metaphorical interpretation is less likely. At home, the metaphorical one might prevail. Underlying it could be a representation like Figure 3. (Notice that in addition to the NAME of the Role (of no computational interest) there is the pointer labeled V/R. This stands for Value/Restriction and provides information about the fillers of a Role. V/Rs must be other Concepts. NAME and V/R are two facets that Roles can be thought of as having.) The common person may well consider him or herself to be different from the animals (witness the creationist-evolutionist debates). Then *John is an animal* would no longer fall into a generality, specificity situation and, although PEOPLE and ANIMALS are in a sibling relationship to one another, the relationship is between members of categories far more abstract than basic level categories. Thus, the incongruity would make the metaphorical interpretation more likely.¹¹

¹¹ I am grateful to Robert Dietz and Loretta Hirsch (personal communication) for the example that led me to this refinement.

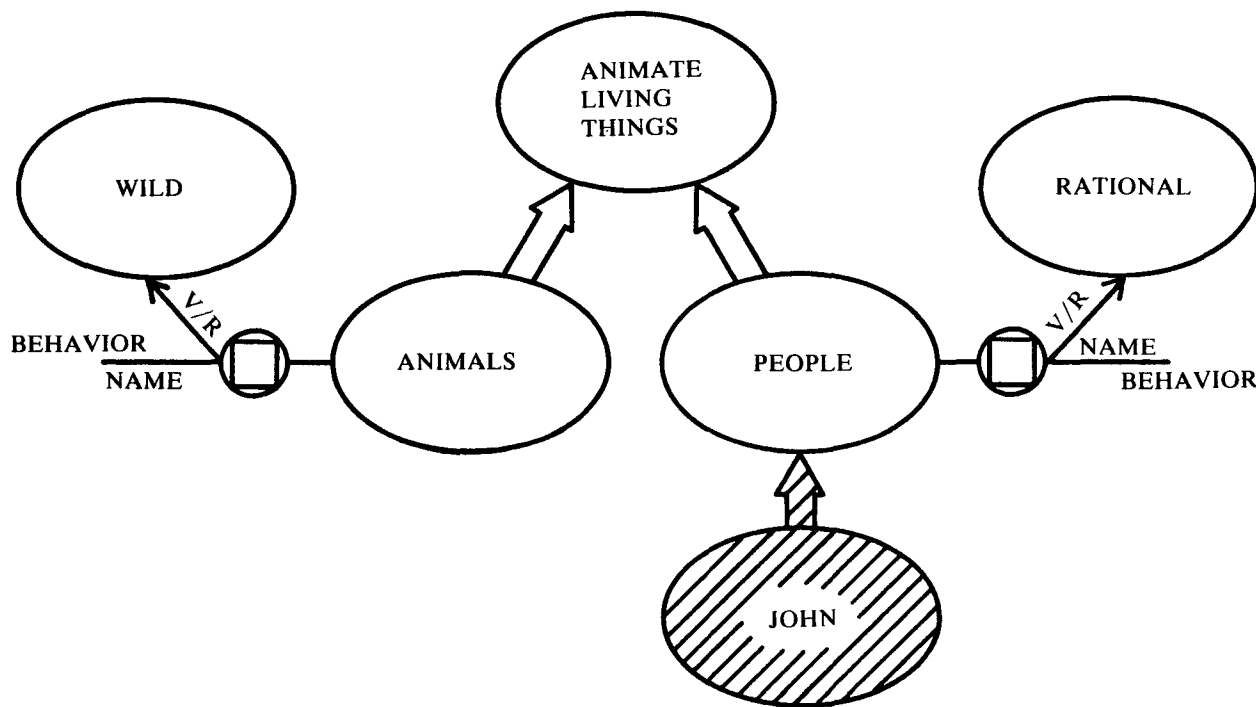


Figure 3.

With metaphors as with other forms of user input, contradictions with the knowledge base may occur. The idea of a user contradicting the knowledge of the system raises the issue of the relationship between these two. The core of knowledge originally stored in the base should be thought of as having been created by an expert to be used by a layman. Consequently, information supplied by the user should have a different status from that of the original designer of the system. However, for communication with the system to take place, discourse elements must at least temporarily be integrated into the knowledge base.¹²

The kind of contradictions most common in metaphors will be those involving a change in salience. For example, the system may know that John is an attractive guy with sloppy eating habits. Both may be considered of equal salience in this case. If the user utters

(27) John is a pig,

then the salience of his eating habits for this user has been elevated beyond that of his otherwise pleasant demeanor. (This is referred to as predicate promotion in Ortony 1979b.) On the other hand,

(28) John is a doll

would have the opposite effect. John's being an attractive guy is a more important characteristic of John for this user. In the somewhat unlikely event that both (27) and (28) were uttered by the same user, the salience of both would be elevated.

How then can a representation system like KL-ONE be used? It should first be noted that in the interest of prototypicality considerations I will follow Cohen (1982b) in allowing V/Rs on Roles of a Concept to include an exclusive disjunction of possible values, weighted by typicality. The lower Concepts can then restrict these to the appropriate ones. These Concepts can in turn also be ordered by typicality. In Figure 4 then, RED is a more typical color for an APPLE than GREEN and a DELICIOUS apple is a more typical apple than a GRANNY SMITH. These rankings are indicated by the symbol ">". This representation of APPLE, having as it does only one Role, is of course highly simplified. Here, color is restricted to RED or GREEN, where RED and GREEN are other Concepts in the network.

If the sentence

(29) Jane's cheeks are like apples

is to be understood, it is necessary to have a Concept of the prototypical apple. Among the members of the community that would understand this, surely the color of prototypical apples is red (see Figure 4 - RED > GREEN and DELICIOUS > GRANNY SMITH).


¹² Others have provided for discourse phenomena in representing knowledge. For example, see sections on semantic knowledge and discourse knowledge in Walker (1978) for a discussion of the discourse component of the SRI speech understanding system. The relationship of context to non-literal language is explored in Ortony, Schallert, Reynolds, and Antos (1978).

(30) Jane's cheeks are like Granny Smith apples
 would certainly have a different interpretation.

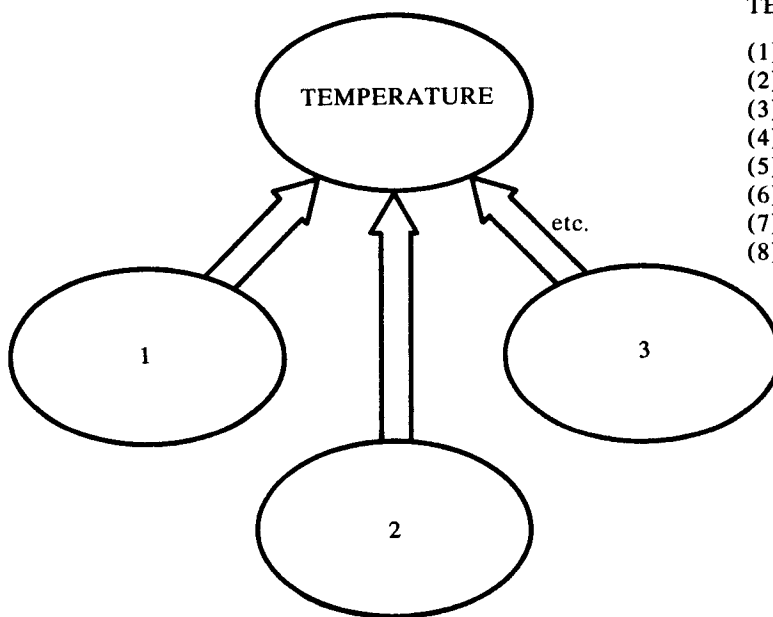
In Figure 5, notice that for the concept HAND, temperature is not terribly salient.¹³ I have also introduced a range of values as a possible value restriction, here 3-6. So this means that hands prototypicality range from hot to cold. In considering the Concept ICE (Figure 6), notice that it has a temperature of 7, that is, extremely cold. Temperature is highly salient for ICE(=1).

Figure 6 also illustrates the relationship between a higher Concept, SOLIDS, and a lower one, ICE. The Roles of SOLIDS are inherited by ICE. In some representations they are inherited intact, but here the Restricts link causes a restriction of the fillers of the Role in question. Thus, TEMPERATURE is restricted to 7; TEXTURE is restricted to HARD/SMOOTH.

¹³ I have provided for salience to assume values between zero (least salient) and one (most salient). The algorithm for the computation of these values awaits further empirical results. At present, they represent an estimate based on my intuitions.



REPRESENTATION OF HAND



REPRESENTATION OF TEMPERATURE

TEMPERATURE:

- (1) BEYOND LINGUISTIC DESCRIPTION (HOT)
- (2) EXTREMELY HOT
- (3) HOT
- (4) LUKEWARM
- (5) COOL
- (6) COLD
- (7) EXTREMELY COLD
- (8) BEYOND LINGUISTIC DESCRIPTION (COLD)

Figure 5.

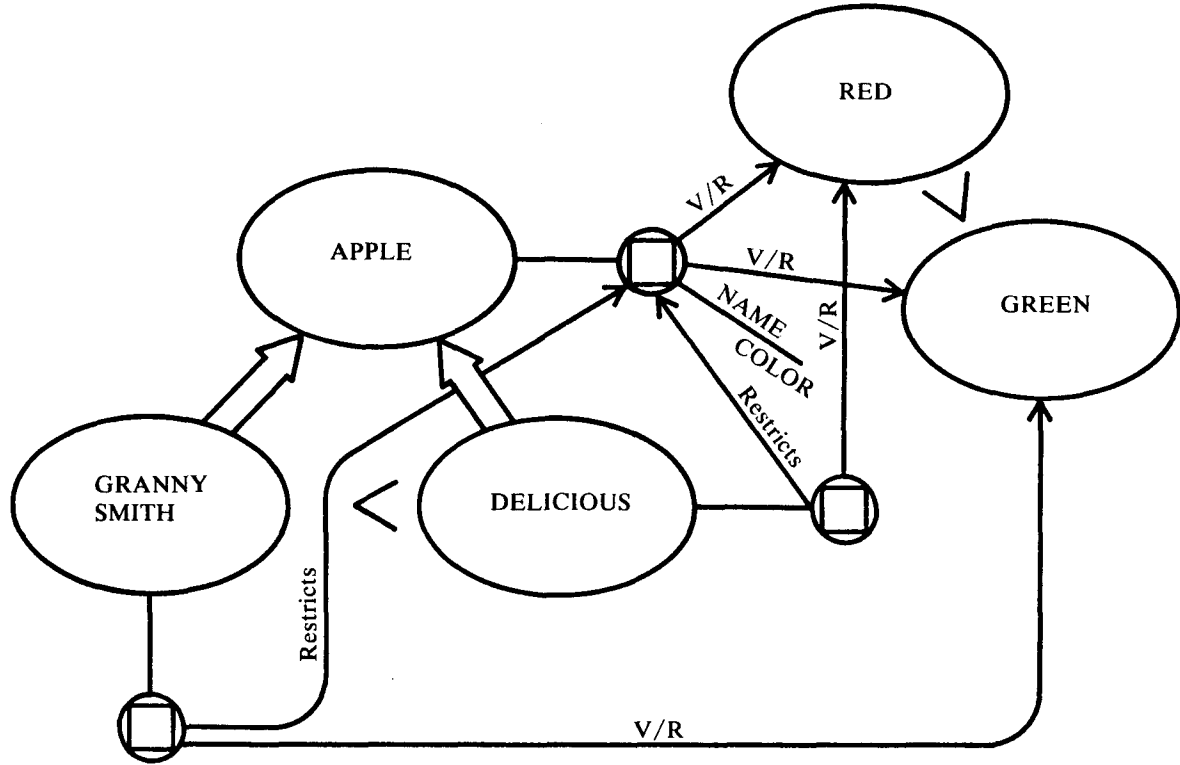


Figure 4.

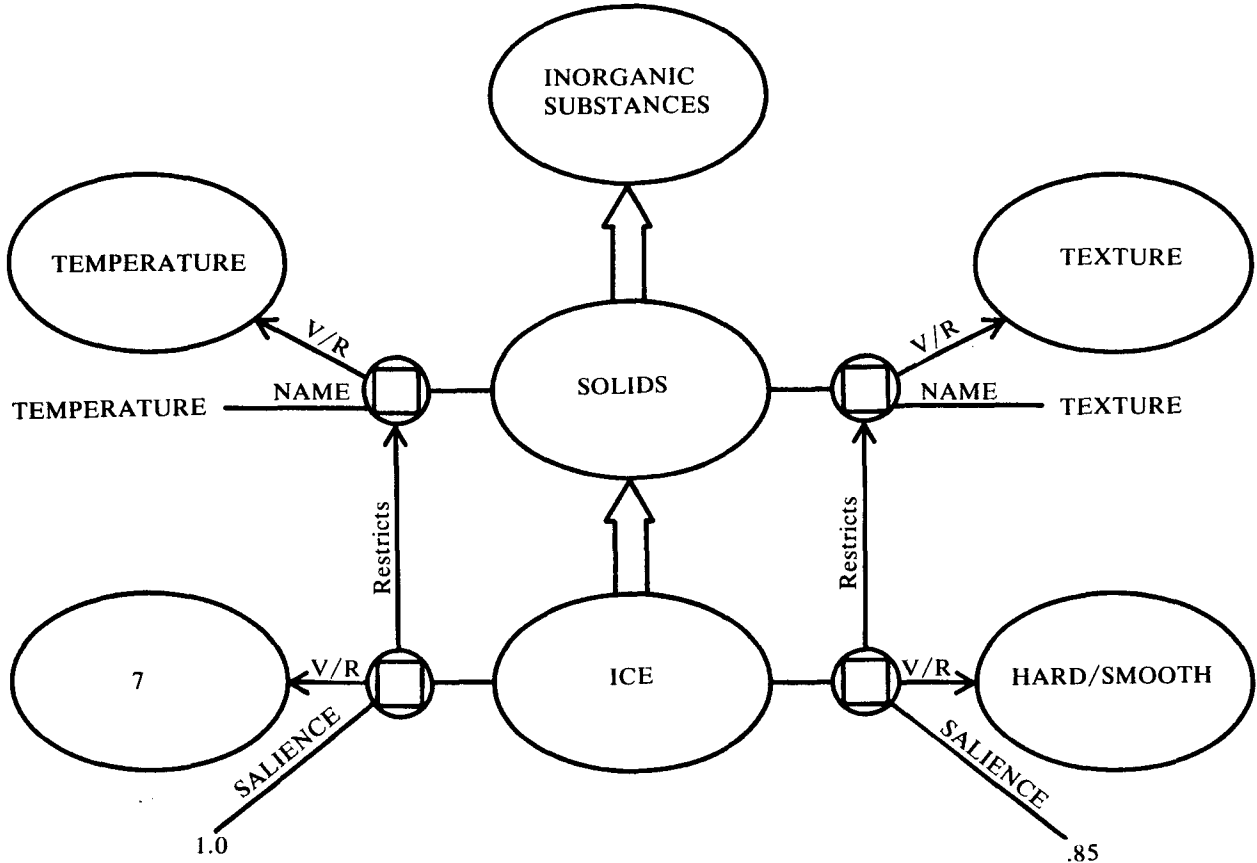


Figure 6.

What then would be the result of the user input

(9) John's hands are like ice.

Looking at a portion of John's network integrated with the one for HAND (Figure 7), sentence (9) would cause one to look at the Roles for John's hands, see that they do not Restrict the Roles of the Generic Concept HAND. Therefore, they are thought by the system to be prototypical hands. Temperature is the most salient feature of ICE for which John has a Role. This must be the Role indicated by (9). The user has made a hyperbolic statement. He or she has said that John's hand temperature is inexpressible by reference to normal hand temperatures. Thus, to understand the sentence, it is necessary to observe that ICE has a much more extreme value for temperature and it is of the highest salience (=1). As a result, for the purposes of this discourse, the salience of John's hand temperature is given the value of 1, implying that for the speaker John's hands are one of his most salient features. This information may be useful in interpreting the discourse that follows the sentence.

The incongruity that must be present for metaphors to work can also be seen by referring to these diagrams, particularly with regard to their hierarchical nature. It is clear that a metaphorical statement is

possible between members of an inheritance relationship

(31) John is a person

(32) Ice is a solid

since these are actual statements of that relationship. A deeper analysis of human classification devices promises to yield further constraints on pairs that can relate metaphorically to one another. For example, looking back at

(12) John is like his father,

since both John and his father share the same immediate superordinate category, male, (12) cannot be metaphoric. This may also explain why

(33) Encyclopedias are like gold mines

is metaphorical but

(34) Encyclopedias are like dictionaries

seems to be a similarity statement. (Examples from Ortony, Reynolds, and Arter 1978.)

The algorithm implied here can be expressed as follows:

1. If the topic is an individual constant (IC), establish restrictions (using the Restricts link), if any, on the Role in question (for example, JOHN's HANDS). If

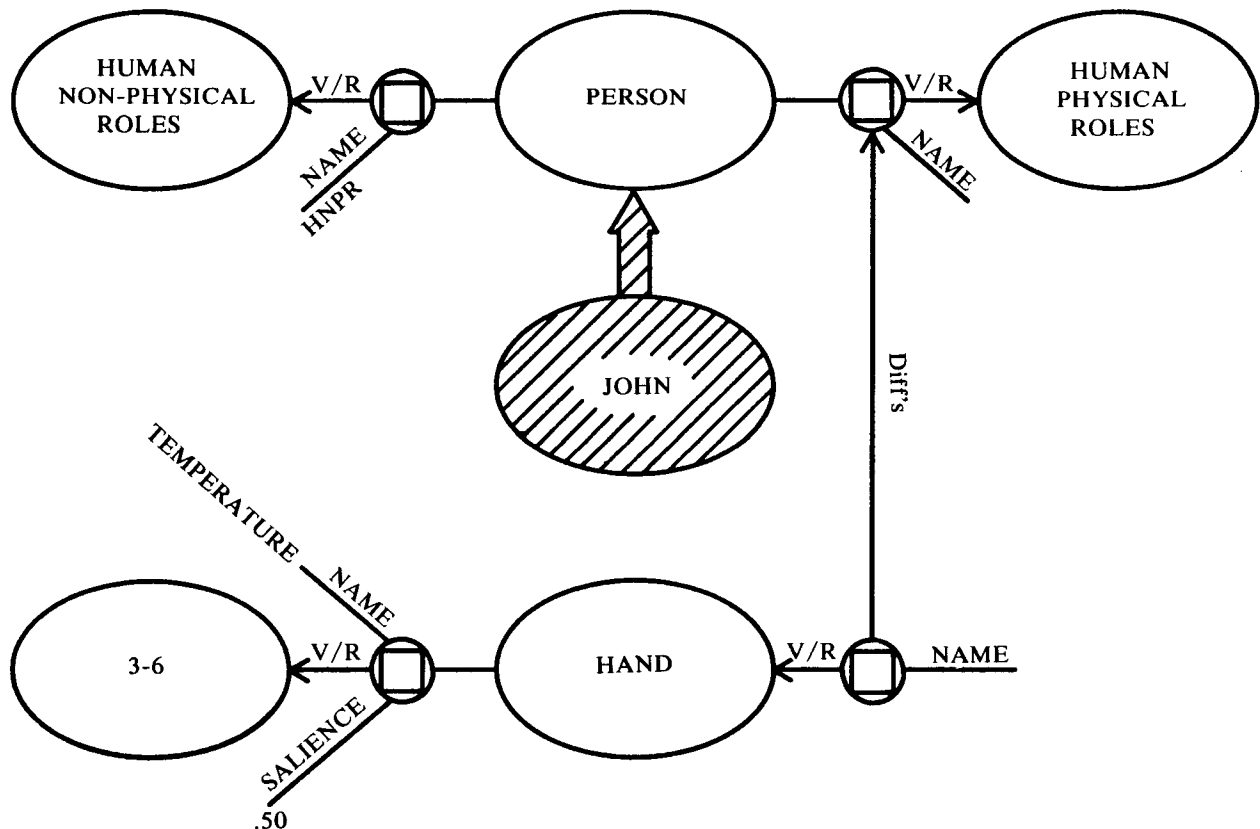


Figure 7.

there are restrictions, note these; otherwise, note inherited V/Rs.

2. Establish those salient predicates for the vehicle for which the topic also has a Role (for example, TEMPERATURE for ICE and JOHN's HANDS).
3. If the V/Rs for these Roles are extreme in the vehicle but not in the topic, the utterance is hyperbolic. If, in addition, the vehicle and topic are in the proper relationship to one another with respect to the taxonomy, the utterance is metaphorical. Given that these conditions hold, raise the salience of the relevant Roles of the topic.

We have begun the computer implementation of this algorithm using NIL on a VAX 11/780. We have implemented enough of the features of KL-ONE to allow us to build a prototype knowledge base (in progress) that will be rich enough to permit experimentation using input consisting of novel metaphors. We intend to exercise the system with the goal of establishing the correctness or need for refinement of the algorithm.

12. Conclusions

To summarize, I have demonstrated how metaphor comprehension can proceed on the foundation of salience, with the following modifications to the theories of Ortony: High salient predicates of the A term can be those at issue in a metaphorical (as opposed to strictly literal) interpretation and low salient features of the B term likewise are of concern. Because of other factors that serve to motivate the use of metaphors (incongruity, hyperbole, inexpressibility), metaphors are not always compact, nor are they prohibited from being used for a single predicate.

Prototype theory applies in two ways: the B term is generally chosen as prototypical of certain predicates; the real-world representative of the B term is a prototypical member of its class.

In my utilization of KL-ONE, I have allowed for a range of possible values in the value restriction facet and introduced salience as a role facet. I have also demonstrated the need for sub-knowledge networks in dealing with metaphors and other natural language issues.

13. The Large Scope of Things

Although this approach to natural language and knowledge representation has been from the point of view of metaphors, it seems clear that at least some of the factors operating in metaphor understanding operate in literal language as well. There is undoubtedly a relationship between salience raising in metaphors and the resultant effect on discourse and focus of attention. In fact, salience raising undoubtedly contributes to focusing. I have only dealt lightly with discourse problems and recognize these as crucial to all language

understanding, whether the language be literal or figurative.

It should be noted that the approach I have taken ignores the possibility of considering metaphors as analogies. Since many are not analogies, I will save that for future work. That they exist is clear:

(35) Giraffes are like skyscrapers

is somewhat more complicated to understand (computationally speaking) than many of the ones I have used in this discussion because it involves relationships among Roles and not Roles simply. (Giraffes are the tallest animals, skyscrapers are the tallest buildings.) Also, clearly, some analogies and some metaphors are instructional (Ortony 1975).

(37) The structure of an atom is like the structure of the solar system.

These appear to be used in building the representation of a new concept (here, atoms). I have dealt here only with representations of existing concepts.¹⁴

In addition, I have chosen to develop a method that could be applied to handling novel metaphors as opposed to those recognized by Lakoff and Johnson (1980) as general metaphors. A system could be made more efficient by utilizing a technique such as Carbonell (1981) has described for recognizing the latter whenever they occur and incorporating my proposals elsewhere.

Finally, as in the case of literal language, it is essential to study recordings of natural speech to see what people *actually* say. In metaphors as elsewhere, there will be many surprises.

Acknowledgments

I'd like to thank Ralph Weischedel, Michael Freeman, and Genevieve Berry-Rogghe for reading and commenting on an earlier draft of this paper.

References

- Anderson, R.C. and Ortony, A. 1975 On Putting Applies into Bottles – A Problem of Polysemy. *Cognitive Psychology* 7: 167-180.
- Brachman, R.J. 1977 A Structural Paradigm for Representing Knowledge. Ph.D. thesis. Harvard University, Cambridge, Massachusetts.
- Brachman, R.J. 1978 Theoretical Studies in Natural Language Understanding. Annual Report, 1 May 1977 to 30 April 1978; BBN Report No. 3888 (September). Bolt Beranek and Newman, Inc., Cambridge, Massachusetts.
- Brachman, R.J. 1979 On the Epistemological Status of Semantic Networks. In: Findler, N.V., Ed., *Associative Networks: Representation and Use of Knowledge by Computers*. Academic Press, New York, New York: 3-50.
- Carbonell, J.G. 1980 Metaphor – a Key to Extensible Semantic Analysis. In *Proceedings of the 18th Annual Meeting for the Association for Computational Linguistics and Parasession on*

¹⁴ Ortony (1979b, p. 199) recognizes the distinction in terms of "predicate promotion" versus "predicate introduction".

- Topics in Interactive Discourse*, June 19-22, 1980. University of Pennsylvania, Philadelphia, Pennsylvania: 17-21.
- Carbonell, J.G. 1982 Metaphor: An Inescapable Phenomenon in Natural Language Comprehension. In: Lehnert, W. and Ringle, U., Eds., *Strategies for Natural Language Processing*. Lawrence Erlbaum Associates, Hillsdale, New Jersey: 415-434.
- Cohen, B. 1982a A Knowledge Representation Model of Prototype Theory. In: Schmolze, J.G. and Brachman, R.J., Eds., *Proceedings of the 1981 KL-ONE Workshop*. BBN Report No. 4842 (June). Bolt Beranek and Newman, Inc., Cambridge, Massachusetts: 124-127.
- Cohen, B. 1982b Understanding Natural Kinds. Ph.D. Thesis. Stanford University, Stanford, California.
- Conklin, E.J. and McDonald, D.D. 1982 Saliency: The Key to the Selection Problem in Natural Language Generation. In *Proceedings of the 20th Annual Meeting of the Association for Computational Linguistics*, 16-18 June 1982. University of Toronto, Toronto, Canada: 129-135.
- Gentner, D. 1980 Studies of Metaphor and Complex Analogies. Paper presented at the APA Symposium on Metaphor as Process.
- Gentner, D. 1981 Are Scientific Analogies Metaphors? In: Miall, D.S., Ed., *Metaphor: Problems and Perspectives*. Harvester Press, Ltd., Brighton, Sussex, England.
- Grosz, B.J. 1977 The Representation and Use of Focus in Dialog Understanding. Ph.D. dissertation. University of California at Berkeley, Berkeley, California.
- Grosz, B.J. 1981 Focusing and Description in Natural Language Dialogues. In: Joshi, A.; Webber, B.; and Sag, I., Eds., *Elements of Discourse Understanding*. Cambridge University Press, New York, New York.
- Haynes, F. 1975 Metaphor as Interactive. *Educational Theory* 25(3): 272-277.
- Labov, W. 1973 The Boundaries of Words and their Meanings. In: Bailey, C.J.N. and Shuy, R.W., Eds., *New Ways of Analyzing Variation in English*. Georgetown University Press, Washington, D.C.: 340-373.
- Lakoff, G. and Johnson, M. 1980 *Metaphors We Live By*. University of Chicago Press, Chicago, Illinois.
- Ortony, A. 1975 Why Metaphors Are Necessary and Not Just Nice. *Educational Theory* 25(1): 45-53.
- Ortony, A. 1979a Beyond Literal Similarity. *Psychological Review* 86(3): 161-179.
- Ortony, A. 1979b The Role of Similarity in Similes and Metaphors. In: Ortony, A., Ed., *Metaphor and Thought*. Cambridge University Press, New York, New York.
- Ortony, A.; Reynolds, R.E.; and Arter, J.A. 1978 Metaphor: Theoretical and Empirical Research. *Psychological Bulletin* (September):919-943.
- Ortony, A.; Schallert, D.L.; Reynolds, R.E.; and Antos, S.J. 1978 Interpreting Metaphors and Idioms: Some Effects of Context on Comprehension. *Journal of Verbal Learning and Verbal Behavior* 17: 465-477.
- Rosch, E.H. 1973 On the Internal Structure of Perceptual and Semantic Categories. In: Moore, T.E., Ed., *Cognitive Development and the Acquisition of Language*. Academic Press, New York, New York: 111-144.
- Rosch, E. and Mervis, C.B. 1975 Family Resemblances: Studies in the Internal Structure of Categories. *Cognitive Psychology* 7: 573-605.
- Rumelhart, D.E. and Ortony, A. 1977 The Representation of Knowledge in Memory. In: Anderson, R.C.; Spiro, R.J.; and Montague, W.E., Eds., *Schooling and the Acquisition of Knowledge*. Lawrence Erlbaum Associates, Hillsdale, New Jersey: 99-135.
- Schmolze, J.G. and Brachman, R.J. 1982 Summary of the KL-ONE Language. In: Schmolze, J.G. and Brachman, R.J., Eds., *Proceedings of the 1981 KL-ONE Workshop*. BBN Report No. 4842 (June). Bolt Beranek and Newman, Inc., Cambridge, Massachusetts: 233-260.
- Tversky, A. 1977 Features of Similarity. *Psychological Review* 84(4) 327-352.
- Verbrugge, R.R. 1980 Transformations in Knowing: A Realist View of Metaphor. In: Honeck, R.P. and Hoffman, R.R., Eds., *Cognition and Figurative Language*. Lawrence Erlbaum Associates, Hillsdale, New Jersey: 87-125.
- Walker, D.E., Ed. 1978 *Understanding Spoken Language*. Elsevier North-Holland, Inc., New York, New York.
- Weiner, E.J. 1982 Using KL-ONE to Understand Metaphors. In: Schmolze, J.G. and Brachman, R.J., Eds., *Proceedings of the 1981 KL-ONE Workshop*. BBN Report No. 4842 (June). Bolt Beranek and Newman, Inc., Cambridge, Massachusetts: 194-197.
- Weiner, E.J. 1983 The Role of the Knowledge Network in the Computational Processing of Non-literal Language. In *Proceedings of the International Conference on Data Bases in the Humanities and Social Sciences*. (in press)
- Williams, O., Ed. 1963 *The Mentor Book of Major British Poets*. The New American Library, Inc., New York, New York.