METAPHOR INTERPRETATION AS SELECTIVE INFERENCING

Jerry R. Hobbs
SRI International
Menlo Park, California

Abstract

Metaphor pervades natural language discourse. This paper describes a computational approach to the interpretation of metaphors. It is based on a natural language processing system that uses the discourse problems posed by a text to select the relevant inferences. The problem of interpreting metaphors can then be translated into the problem of selecting the relevant inferences to draw from the metaphorical expression. Thus, a metaphor is frequently given a correct interpretation as a by-product of the other things a natural language system has to do. Two examples of metaphors are given — a spatial metaphor schema from computer science, and a novel metaphor — and it is shown how the interpretation problem for each can be translated into a selective Inferencing problem and solved by the ordinary operations of the system. This framework sheds light on the analogical processes that underlie metaphor and begins to explain the power of metaphor.

1. Metaphor is Pervasive

I. A. Richards, in speaking of metaphor, said, "Literal language is rare outside the central parts of the sciences." (Richards 1936). But it is rare even in the central parts of the sciences. Consider for example the following text from computer science. It comes from an algortihm description in the first volume of Knuth's <u>Art of Computer Programming</u>, Vol. 1, p. 417, and is at but one remove from the domain's most formal mode of expression.

"Given a pointer PO, this algorithm sets the MARK field to 1 in NODE(PO) and in every other node which can be reached from NODE(PO) by a chain of AL1NK and BLINK pointers in nodes with ATOM - MARK = 0. The algorithm uses three pointer variables, T, Q, and P, and modifies the links and control bits during its execution in such a way that all ATOM, ALINK, and BLINK fields are restored to their original settings after completion, although they may be changed temporarily."

In this text the algorithm, or the processor that executes it, is apparently a purposive agent that can perform such actions as receiving pointers; setting, changing, and restoring fields; reaching nodes; using variables for some purpose; modifying links and bits; and executing and completing its task.

Nodes are apparently locations that can be linked and strung into paths by pointers and visited by the processor-agent.

Nodes also seem to be containers that can contain fields.

Fields are also containers that can contain pointers, among other things. In addition, fields are entitles that can be placed at, or set to, locations on the number scale or in the structured

their very name, suggest objects that can point to a location for some agent's information.

In fact, there is very little in the paragraph that does not rest on <u>some</u> spatial or agent metaphor. Moreover, these are not simple isolated metaphors; they are examples of large-scale "metaphor schemata" or "root metaphors" (Lakoff and Johnson 1980) which we use to encode and organize our knowledge about the objects of computer science. They are so deeply ingrained that their metaphorical character generally escapes our notice.

The pervasiveness of metaphor was noted as early as the eighteenth century by Giambattista Vico (1744, 1968). In our century, this observation has been the basis for a rejection of Aristotle's and Quintillian's views that metaphor is mere ornament, and an elevation of metaphor to an "omnipresent principle of language" (Richards 1936) and "the law of its life" (Langer 1942). As we saw in our example, the spatial metaphor especially is pervasive. Jespersen (1922) remarked on this. For Whorf (1939) it was a key element in his view that language determines thought: the spatial metaphors provided by one's langauge determines how one will normally conceptualize abstract domains. The most thoroughgoing recent treatment of metaphor in everyday language is found in Lakoff and Johnson (1980); they identify the core metaphors that underlie our thinking about a vast array of domains, and argue that we can understand the domains only by means of these metaphors. The Fundamental Insight that informs all this work is this: metaphor is pervasive in everyday discourse and is essential in our conceptualizations of abstract domains. The aim of this paper is to present a computational treatment of metaphor that can accommodate this Fundamental Insight.

2. Some Frameworks for Investigating Metaphor

The earliest detailed proposal in computational linguistics for handling metaphor was that of Russell (1976). Her proposal concerns abstract uses of verbs of motion and involves lifting selectional constraints on the arguments of the verb while keeping fixed the topological properties of the motion, such as source, path and goal. Thus, to handle "the ship plowed through the sea," we lift the restriction on "plow" that the medium be earth and keep the property that the motion is in a substantially straight line through some medium.

Russel exemplifies an approach that finds its most complete development in Levin (1977), Metaphor is treated as a species of semantic deviance; selectional constraints are lifted until the expression can sail through the interpreter without difficulty and without effect. But the problem of interpreting "the ship plowed through the sea" is not to avoid rejecting the sentence because the sea is not earth, but to notice the similarity of the wedge-shaped plow and the wedge-shaped bow of a ship and the wake that each leaves, and perhaps more importantly, to take note of the ship's steady, inexorable progress. Any approach to metaphor that does only the first of these is not a way of interpreting metaphors, only of Ignoring them. Under this view, the Fundamental Insight is simply bizarre and inexplicable.

A more productive approach to metaphor interpretation can be based on work in mainstream natural language processing. One of the principal thrusts of natural language processing research in the last decade has been to develop frameworks in which inferences can be drawn selectively. One reason that such systems are needed is that it is difficult, if not impossible, to axiomatize in a consistent manner any domain more complex than set theory. However another reason is that there are too many true inferences that can be drawn in a specific situation and most of them are irrelevant. A great deal of work in natural language processing can be viewed as addressing this problem (including Hobbs 1976, Joshi and Rosenscheln 1976, Grosz 1977, Schank and Abelson 1977, and Mann, Moore and Levin 1977).

Metaphor interpretation in such a framework becomes the problem of drawing certain inferences and refraining from drawing other inferences. Consider a simple case. Suppose it were not a cliche to call someone a hog. How would we go about interpreting the sentence

(1) John is a hog.

Let us suppose our initial logical representation for this is

hog(J).

There are a number of things we might infer from the fact that some entity is a hog, among them

hog(x) —> fat(x)
hog(x) —> overcon8ume(x,y),food(y)
hog(x) —> sloppy(x)
hog(x) —> has-four-legs(x)

The problem we are faced with in interpreting (1) is the problem we are always faced with in interpreting a text — determining which inferences it is appropriate to draw from what we've been told. Depending on the situation, we may want to infer "fat(J)", or "overconsume(J,F)" where "food(F)", or simply "overconeume(J,X)" where X is some other quantity such as a road John is driving on. The inference that John has four legs is presumably rejected because of strong reasons to believe the contrary. One may or may not infer that John is sloppy, depending on context or other factors.

Our approach then is to say that "John is a hog" conveys at least the information that "John is a hog", but that the various inferences that one could draw from the sentence,

Porky is a hog,

are simply not available to us in the case of John. In particular, most of the inferences that correspond to the various features of the visual image evoked by the word "hog" are not appropriate.

If interpreting a metaphor is a matter of selecting the right inferences, then we must ask how the inferences are to be selected. There have been several interesting proposals.

One is the proposal of Ortony, whose notion of "feature" is probably equivalent to the AI notion of inference. In a metaphor one is comparing an entity in a "new domain" (Richards (1936) called this entity the "tenor", Ortony the "topic") with an entity in an "old domain" (the "vehicle"). Ortony (1979) has suggested a breakdown of the knowledge in the old and new domains into classificatory facts, other high-salience facts, and low-salience facts. Classificatory facts are not transferred from the vehicle to the tenor. Thus, from "John is a hog" we do not infer that John is a farm animal. What get transferred from the vehicle to the tenor are other high-salience facts whose correlates in the tenor are of low salience. It is a high-salience fact that hogs overconsume, a low-salience fact that John overconsumes. The effect of the metaphor is to bring to the fore this low-salience fact about John.

Another interesting proposal is that of Winston (1978). He presents an algorithm in which properties are transferred from the vehicle to the tenor if they are extremes on some scale, are known to be important, or serve to distinguish the vehicle from other members of its class. Thus, properties of hogs that were not shared by other farm animals would be transferred.

Recently Carbone11 (1981) has proposed a division of inferences into a helrarchy with inferences about goals, planning and causes at the top and monadic descriptive inferences and inferences about relevant objects at the bottom. He argues that those near the top of the hierarchy are carried over from the vehicle to the tenor more often than those near the bottom. Gentner (1980) has similarly shown that relations, i.e. dyadic predicates, are carried over more often than monadic predicates.

However, none of these approaches takes into account the text in which the metaphor is embedded. The approach taken in this paper is to subsume the metaphor interpretation problem under the more general problem of making sense of a discourse as a whole. In previous work (e.g. Hobbs 1976) I have investigated the idea that the inferences that it is relevant to draw are the inferences required to solve various discourse problems, like recognizing coherence, forcing congruence between predicates and their arguments, and anaphora and ambiguity resolution. To take a simple (literal) example, consider

(2) John picked up a book.

It la sometimes true that books have tables of contents, and sometimes It is relevant. But there is no reason, given (2) alone, that we would necessarily want to draw the Inference that John'a book has a table of contents. However, if the next sentence in the text is

He turned to the table of contents,

then we can be sure that the inference is both true and relevant.*

To take a metaphorical example, consider

Mary eats like a bird, but John is a hog.

In order to recognise the contrast coherence relation (see Hobbs 1978) indicated by "but", we must draw the inference that John overconsumes food. Other possible inferences about hogs are not drawn, not so much because they would result in an inconsistency, but because no discourse problem requires them to be drawn.

More specifically, I am assuming a framework exemplified by the DIANA system, implemented at SRI for the analysis of discourse. The system may be described briefly as follows. It accepts a text translated by a syntactic front-end into predicate calculus formulae, and draws those inferences necessary to solve the discourse problems posed by the text. The Inferenclng process is selective and driven by a collection of discourse operations which try to do such things as resolve pronoun and definite noun phraae references, find the specific interpretations of general predicates in context ("predicate interpretation"), reconstruct the implicit relation between the nouns in compound nomlnals, and recognise coherence relations between successive portions of the text. The operations select inferences from a large collection of axioms representing knowledge of the world and the language. Associated with the potential inferences are measures of salience which change as the context changes. These help determine which inferences are drawn by the operationa and hence how the text la interpreted. The control structure is such that the system does not try to solve the discourse problems independently, but rather seeks the moat economical interpretation of the sentence as a whole. A more thorough description of the principles underlying this system can be found in Hobbs (1976, 1977, 1979, 1980s) and ita output la exhibited in Hobbs (1980b).

Next we go through two examples in detail — a spatial metaphor schema and a novel metaphor. The first haa been handled by the DIANA system. The second has not, but out framework nevertheless sheds light on how the interpretation would proceed.

Note that we still need the normative knowledge that books often have tables of contents, even though the text is quite explicit. If John had turned to the door, we would not have assumed the book had a door.

3. A Spatial Metphor Schema

Metaphors that tap into our spatial knowledge are especially powerful since our knowledge of spatial relationships is so extensive, so rich, and so heavily used. As soon aa the basis for the spatial metaphor is established, then in our thinking about a new domain we can begin to borrow the extensive machinery we have for reasoning about spatial relationships. For example, once I say that

(3) N is at sero,

and interpret it as

(4) The value of N is equal to sero,

then I have tapped into a large network of other possible uses. I can now say

N goea from 1 to 100

to mean

The value of N successively equals integers from 1 to 100.

I can say

N approaches 100

to mean

The difference between 100 and the value of N becomes smaller.

the simple Identification of (3) and (4) we have bought into the whole complex of spatial terminology.

In terms of a system for selective Inferenclng, what we mean when we aay that our spatial terminology la an intricate network la that there are a great many axioms that relate the various apatlal predicates. The concept of location — the predicate "at" — la at the heart of this network because ao many of the axioms refer to it. For example, we might define "go" by means of axioms like

go(x,y,s) & at(wl,x,y) * at(w2,x,s)
—> become(wl,w2)

that is, if x goes from y to s and w1 is the condition of x being at y and w2 la the condition of x being at z, then there la a change of state, or a "becoming", from wl to w2. Similarly, part of the meaning of "switch" can be encoded in the axiom

swltch(x,yl,y2) & at(wll,yl,sl) 6
at(wl2,yl,s2) 6 at(w21,y2,zl)
* at(w22,y2,z2) —> cauee(x»
become(and(wll,w22),and(vl2,w21))).

That is, if x switches y1 and y2 and wij is the condition of yl being *t sj, then x causes a "becoming" or transition from the state in which wll and w22 hold to a state in which wl2 and w21 hold.

We vert able to establish the metaphor "a variable at an entity at a location" simply by identifying (3) and (4). In our formalism we can eatablish the metaphor with similar simplicity by including the following axiom:

(5) variable (x) * value(w,y,x) —> at(v,x,y)

That is, if x is a variable and w la the condition of y being its value, then w is also the condition of x being at y.

This simple device of identifying "equality" with "being at" gives ua entry into an entire metaphor schema. The schema la repreaented by a collection of axioms that are intricately woven together by their reference to a small aet of common predicates. The schema la tapped for metaphorical purposes by means of axioms like (5), enabling us to transfer to one domain the structure of another, more thoroughly understood domain.

A discourse operation, which in Hobba (1977) waa called predicate interpretation, uses axioms like (S) to arrive at interpretations of certain metaphorical expressions. The idea behind the operation is that most utterances make very general or ambiguous sorts of predications and that part of the job of comprehension is to determine the very specific or unambiguous meaning that was intended. Thus, someone might make the general statement,

I went to London,

expecting us to be able to Interpret it as

1 flew to London in an airplane,

rather than as swimming, sailing, walking, or any of the myriad other manners of going. In the case of (3), we are expected to determine which of the many ways one thing can be at another is intended in this particular caae. That is, rather than determining what we can infer from what is said, we try to determine what the speaker had in mind from which he inferred what he said. In terms of our notation, suppose C is a general proposition and S a specific one and

S —> C

is an axiom expressing a fact that a speaker and a listener mutually know. The speaker utters G in the expectation that the listener will interpret t as S. The listener must locate and use the axiom to determine the specific iinterpretation.

In this manner, axiom (5) provides one possible interpretation of (3), in that it specifies one of the many ways in which one thing can be at another, which the speaker may have meant. When a metaphorical use of "go" or "switch" or any of the other spatial predicates is encountered, axiom (5) combines with the axioms defining the spatial predicate in terms of "at" to give us the correct interpretation.

An alternative to this approach might seem to be to infer intended meaning from what was said. We would use axioms not of the form "3—>G" but of the form

G & Cl & ... & Cn --> M

where G is the general proposition that is explicitly conveyed, the Ci's are conditions determinable from context, and M is the intended meaning. For interpreting (3), this would require an axiom like

(6) at(w,x,y) * variable(x) --> value(w,y,x),

that is, if w is the condition of x being at y and x ia a variable, then w is also the condition of y being the value of x. To interpret (3) we would search through all axioms for axioms that, like (6), have "at" in the antecedent, check whether the other conjuncts in the antecedent were true, and if so, conclude that the axiom's consequent was the intended meaning. This would be equivalent to a "discrimination-net" approach to word-sense disambiguation (e.g. Rieger 1978), in which one travels down a tree-like structure, branching one way or the other according to whether some condition holds, until arriving at a unique specific interpretation at the bottom. The difficulty with this-approach is that it supposes we could anticipate at the outset all the ways the meaning of a word could be influenced by context. For metaphors we would have to be able to decide beforehand on all the precise conditions leading to eech interpretation. It is highly implausible that we could do this for familiar metaphors, and for novel metaphors the whole approach collapses.

4. A Novel Metaphor

Our next example Illustrates how we can repreaent a metaphor that depends on an elaborate analogy between two complex processes. The metaphor comes from a Newsweek article (July 7, 1975) about Gerald Ford's vetoes of bills Congress has passed. A Democratic congressman complains:

(7) We insist on serving up these veto pitches that come over the plate the else of a pumpkin.

It is clear from the rest of the article in which this appears that this means that Congress has been passing bills that the President can easily veto without political damage. There are a number of problems ralaed by this example, but the only onea we will addrees are the questions of how to represent and Interpret "veto pitches that come over the plate".

The analogy here is between Congress sending a bill to the Preeldent to sign or veto end a pitcher throwing a baseball past a batter to miss or hit. Let us try to encode each of the processes first, then establish the links between them, and then show how a natural language system might discover them.

The facta about a bill are as follows. The participants are Congress, the bill, and the President. Congress sends a bill to the President, who then either signs it or vetoes it. We assume there is an entity C, Congress. To encode the fact that C is Congress, we could write

Congress(C).

But it will prove more useful to assume there is a condition — call it CC — which is the condition of C being Congress. We represent this

Congress(CC,C).

Similarly, there are entities B, CB, P, and CP, with the properties

bIll(CB,B),

i.e. CB is the condition of B being a bill, and

Preeident(CP,P),

i.e. CP is the condition of P being the President. There are three relevant actions, call them SD, SG, and VT_V with the following properties:

send(SD,C,B,P),

i.e. SD is the action by Congress C of sending the bill B to the President P;

sign(SG,P,B),

i.e. SG is the action by the President P of signing the bill B; and

veto(VT,P,B),

i.e. VT is the action by the President P of vetoing the bill B. There is the condition — call it OSV — in which either the signing SG takes place or the vetoing VT takes place:

or(OSV,SG,VT).

Finally, there is the situation or condition, TH, of the occurrence of the sending SD followed by the alternative actions OSV:

then(TH,SD,OSV).

The corresponding facts about baseball are as follows: There are a pitcher x, a ball y, and a batter z, and there are the conditions ex, cy, and cz, of x, y, and z being what they are:

pItcher(cx,x)
ball(cy.y)
batter(cs,z).

The actions are the pitching p by the pitcher x of the ball y to the batter s,

pitch(p,x,y,z);

the missing m of the ball y by the batter z,

mIss(a,s,y);

and the hitting h of y by z,

hit(h,y,z).

Let omh represent the condition of one or the other of m and h occurring,

or(omh,m,h),

and th the situation of the pitching p followed by either m or h,

then(th,p,omh).

Where individual constants, C, CC, B, ..., were used in the Congressional bill schema, universally quantified variables, x, ex, y, •••, are used here. This is because the baseball schema is general knowledge that will be applied to the specific situation involving Congress and the President. It is a collection of axioms that get instantiated in the course of interpreting the metaphor.

The linkage established by the metaphor is among other things, between the bill and the ball. But it is not enough to say that B, in addition to being the bill, is also in some sense a ball, just as B has other properties, say, being concerned with federal housing loans, being printed on paper, and containing seventeen subsections. The metaphor is stronger. What the metaphor tells us is that the condition of B being the bill is indeed the condition of B being a ball. Similar links are established among the other participants, actions, and situations. That is, the baseball schema is instantiated with the entitles of the Congressional bill schema, leading to the following set of propositions:

(8)

Congress(CC,C)	**pitcher(CC,C)**
bill(CB,B)	**ball(CB,B)**
President(CP,P)	**batter(CP,P)**
send(SD,C,B,P)	**pitch(SD,C,B,P)**
sign(SG,P,B)	**miss(SG,P,B)**
veto(VT,P,B)	**hit(VT,P,B)**

or(OSV,SG,VT)
then(TH,SD,OSV)

Although all of this has been described in terms of schemata, a schema in this framework is simply a collection of possibly very complex axioms that are interrelated by the co-occurrence of some of the same predicates, perhaps together with some metaknowledge for controlling the use of the axioms in InferencIng. The linkage between the two schemata does not require some special "schema-mapping" operation, but only the assumption of identity between the corresponding conditions, just as in the previous example we identified "equality" with "being at". It is because of the mechanisms of selective InferencIng that this will do. Thus, to represent the metaphor, we do not have to extend our formalism beyond what was required for the first example, nor indeed beyond what is required for nonmetaphorIcal discourse.

No natural language processing system existing today could derive (7) from (8). Nevertheless, we can see the basic outline of a solution within the selective InferencIng framework: The congressman said, "We insist on serving up these veto pitches ••••" For someone to serve up a pitch is for him to pitch. This leads to the identification of

Congress with the pitcher. To Interpret the compound nominal "veto pitch", we must find the most salient, plausible relation between a veto and a pitch* From our knowledge about vetoes, we know that Congress auet first **send** the bill to the President. Proa our knowledge about pitching, we know that for the Congress/pitcher to pitch, It must **send** a "ball" to a "batter". We have a match on the predicate "send" and on the agents of the sendlngs, Congress. We can complete this match by assuming the bill is the ball and the President is the batter.

Ve have almost a complete match between the two situations. The analogy will be completed when we determine which of the various possible actions that a batter can perform corresponds to the President's veto. But this is just what we need to complete the relation between "veto" and "pitch" in the compound nominal. By some means well beyond the scope of this paper to discuss, "pitches that come over the plate the site of a pumpkin" must be interpreted to mean that the ball is easy for the batter to hit. If we assume maximum redundancy — that a veto pitch and a pitch that comes over the plate the else of a pumpkin are roughly the same thing — then we assume that the pitch is a bill/ball that the Congress/pitcher sends to the President/batter which he then finds eaay to veto/hit. The analogy is complete.

As with all metaphorical expressions, as indeed with any expression, there will be a number of inferences we will not want to draw in this case — for example, that B is spherical and has stitching. But this metaphor invokes other inferences that we do accept, inferences that would not necessarily follow from the facts about the American government. It suggests, for example, that Congress and the President are adversaries in the same way that a pitcher and a batter are, and that from the President's perspective it is good for him to veto a bill Congress has passed and bad for him to sign it. What we know about the adversary relationship in baseball is vivid and unambiguous, and herein lies the power of the metaphor.

5. Metaphor and Analogy

In the examples of Sections 3 and 4, aa well as in the "John is a hog" example of Section 2, we have seen the same broad processes at work. They can be summarised as follows: There are two domains, which we may call the <u>new domain</u>, or the domain which we are seeking to understand or explicate, and the **<u>old domain</u>**, or the domain in terms of which we are trying to understand the new domain and which provides the metaphor. Richards (1936) refers to these as the tenor end the vehicle, respectively. In our examples the new domains are John's nature, computer science, and the workings of the American government. The old domains are a hog's nature, spatial relationships, and baseball. For each old domain, we can distinguish between what may be called the <u>basic concepts and relationships</u> and <u>complex concepts and relationships</u>. For spatial relationships, "at" is a basic concept; "go", "approach", and "switch" are complex concepts. For baseball, "pitcher" and "batter" are basic, their adversary relationship is complex. What is basic and what is complex in a particular domain are not necessarily fixed beforehand, but may be determined in part by the metaphor itself.

Bach of the examples can be viewed as setting up a link between the basic concepts of a new domain and an old domain, in order that complex concepts or relationships will carry over from the old to the new. The following diagram illustrates this:

```
Old Domain                      New Domain

Complex concepts   -------> Complex concepts
and relationships           and relationships
       A                           A
       |                           |
       |                           |
Basic concepts     <------- Basic concepts
and relationships           and relationships
```

Figure 1.

This diagram is familiar from Galois theory, algebraic topology, and category theory (e.g. Artin 1959). One can prove theorems in one domain (for example, the category of fields) by constructing a "functor" to map its objects and relations into the objects and relations of another domain (for example, the category of groups), proving the theorem in the second domain, and using the inverse functor to map It back into the original domain.

The diagram illustrates a general paradigm for analogical reasoning. To reaaon in a new domain about which we may know little, we map it into an old domain, do the reasoning in the old domain, and map the results back into the new domain. To make use of this paradigm, in our framework, for understanding the proceases of metaphor, we have had to specify the nature of the links in the diagram. The horisontal links are realised by means of axioms like (5) in the case of frozen metaphors and by means of impllcatures like (8) in the case of novel metaphors. The vertical links in the diagram are realised by the collections of axioms encoding the relationships between basic and complex concepts.

in this framework, we can begin to understand why metaphors are used and why they are so pervasive. Any discourse is built on a shared knowledge base of possible inferences. By means of his utterances, the speaker plants Inferences in the listener's head. The richer the shared knowledge base, the more economical, or equlvalently, the more suggestive, the discourse can be. Metaphor la a deceptively simple device for enlarging the knowledge base. By using an apt metaphor to map a new, uncertainly understood domain into an old, well-understood domain, such as

Such assumptions are common in interpreting discourse. In fact, they constitute one of the principal mechanleme for resolving pronouns snd implicit arguments (see Hobbs 1979).

spatial relationships, we gain access to a more extensive collection of axioms connecting the basic and complex levels, thereby securing a more certain grasp on the new domain conceptually and providing it with a richer vocabulary linguistically. It is reasonable to hope that much of the tine the ordinary discourse operations, based on selective inferencing, will insure that the right inferences are drawn and the wrong ones aren't.

Acknowledgments

I have profited from discussions with Armar Archbold, Jaime Carbonell, Gary Hendrlx, George Lakoff, and Bob Moore, and also with the participants of the Artificial Intelligence Workshop at the Electrotechnical Laboratory, Tokyo, Japan, in August 1979. They are, of course, not responsible for any errors in this paper. This work was supported by the National Science Foundation under Grant No. MCS-78-07121 and by the Advanced Research Projects Agency under Contract No. N00039-79-C-0118 with the Naval Electronics System Command.

REFERENCES

Artin, E. 1959. Galois theory. Notre Dame Mathematical Lectures, No. 2. Notre Dame, Indiana.

Carbonell, J. 1981. Metaphor: An inescapable phenomenon in natural language comprehension. Manuscript. May 12, 1981.

Grosz, B. 1977. The representation and use of focus in dialogue understanding. Stanford Research Institute Technical Note 151, Stanford Research Institute, Menlo Park, CA. July 1977.

Hobbs, J. 1976. A computational approach to discourse analysis. Research Report 76-2, Department of Computer Sciences, City College, City University of New York. December 1976.

Hobbs, J. 1977. Coherence and interpretation in English texts. Proceedings, International Joint Conference on Artificial Intelligence, 110-116. Cambridge, Mass. August 1977.

Hobbs, J. 1978. Why is discourse coherent? SRI Technical Note 176, SRI International, Menlo Park, California. November 1978.

Hobbs, J. 1979. Coherence and coreference. Cognitive Science, Vol. 3, No. 1, 67-90.

Hobbs, J. 1980a. Selective inferencing. Proceedings, Third National Conference of Canadian Society for Computational Studies of Intelligence. 101-114. Victoria, British Columbia. May 1980.

Hobbs, J. 1980b. Interpreting natural language discourse. Final Report, National Science Foundation Research Grant No. MCS 78-07121. July 1980.

Jespersen, O. 1922. Language: Its nature, development and origin. London: George Allen & Unwin Ltd.

Joshi, A., and S. Rosenschein. 1976. Some problems of inferencing: Relation of inferencing to decomposition of predicates. Statistical Methods in Linguistics, Vol. 1, No. 1.

Knuth, D. 1973. The art of computer programming. Vol.1. Reading, Mass.: Addison-Wesley.

Lakoff, G. and M. Johnson 1980. Metaphors we live by. Chicago: University of Chicago Press.

Langer, S. 1942. Philosophy in a new key: A study in the symbolism of reason, rite, and art. Cambridge, Mass.: Harvard University Press.

Levin, S. 1977. The semantics of metaphor. Baltimore: The Johns Hopkins University Press.

Mann, W., J. Moore and J. Levin 1977. A comprehension model for human dialogue. Proceedings, International Joint Conference on Artificial Intelligence, 77-87, Cambridge, Mass. August 1977.

Ortony, A. 1979. The role of similarity in similes and metaphors. In A. Ortony (Ed), Metaphor and Thought. Cambridge, England: Cambridge University Press.

Richards, I. A. 1936. The Philosophy of Rhetoric. Oxford: Oxford University Press.

Rieger, C. 1978. Viewing parsing as word sense discrimination. In W. Dingwall (Ed.), A Survey of Linguistic Science. Greylock Pub.

Russell, S.W. 1976. Computer understanding of metaphorically used verbs. American Journal of Computational Linguistics, microfiche 44.

Schank, R. and R. Abelson 1977. Scripts, plans, goals, and understanding. Hillsdale N.J.: Laurence Erlbaum Associates.

Vico, G. 1744. The new science of Giambattista Vico. T. Bergin and M. Frisch, trans. Ithaca NY: Cornell University Press. 1968.

Whorf, B. 1939. The relation of habitual thought and behavior to language. In J. Carroll (Ed.) Language, thought, and reality: Selected writings of Benjamin Lee Whorf, 134-159. New York: John Wiley & Sons. 1956.

Winston, P. 1978. Learning by creating and justifying transfer frames. Artificial Intelligence, Vol. 10, 147-172.