Metaphor and Communication

Boaz Keysar; Sam Glucksberg

Poetics Today, Vol. 13, No. 4, Aspects of Metaphor Comprehension (Winter, 1992), 633-658.

Metaphor and Communication

Boaz Keysar
Psychology, Chicago and

Sam Glucksberg
Psychology, Princeton

Abstract To what extent do communicative principles constrain the compre-
hension of simple metaphors? We argue that the comprehension of metaphor
depends on communicative principles in exactly the same way that literal
comprehension depends on them. We support our claims with our experimen-
tal findings that indicate equivalent processing for literal and metaphorical
language. We then outline our theory of metaphor comprehension, which ac-
counts for a variety of metaphoric phenomena and explains why people use
metaphors.

People can reliably distinguish between literal and figurative language
and can easily identify a literal use of an expression, as opposed to
idiomatic, metaphorical, ironic, and other figurative uses. How is this
done? Standard theories of language suggest that people use different
processes when they interpret an expression as literal or as figurative
and that this difference is the basis for the distinction. In this paper
we will focus on the literal/metaphorical distinction and will start by
arguing against the notion that the interpretative processes of meta-

phor differ from those of literal language. But first we will consider
the role of communication principles in the understanding of literal
and metaphorical expressions.

One of the major unresolved questions in the study of language is
to what extent communicative principles guide the understanding of
language. This question will be present in the background throughout
our discussion, but we will not attempt to answer it. We would, how-
ever, challenge the common belief that communicative principles are
more crucial to the understanding of non-literal than of literal lan-
guage. We suggest that, to the extent that communicative principles
are crucial for metaphor comprehension, they are equally important
for literal-language understanding.

First, we will discuss the general question: Should principles of com-
munication be part of a theory of language understanding? We will
consider two approaches that provide contradictory answers to this
question. One approach assumes that communication is irrelevant to a
theory of language, the other that communication is crucial for under-
standing language. We will then show that despite this disagreement,
these two views share a number of basic assumptions. Both approaches
assume that the understanding of literal language differs in important
respects from the way in which metaphorical language is understood.
The implicit area of consensus between these diametrically opposed
approaches is their agreement that communicative principles under-
lie the understanding of figurative, but not of literal, language. We
will argue against this pervasive idea and suggest that the same basic
operations underlie people's understanding of both literal and figu-
rative language. On the one hand, we will show that those aspects
of comprehension that are not constrained by communication prin-
ciples are shared by literal and metaphorical language; on the other
hand, we will argue that, when communicative principles are used,
they are used similarly in both metaphorical and literal language. We
will then propose a model of metaphor understanding and conclude
with a discussion of the different communicative *functions* that literal
and metaphorical language may serve.

Language and Communication: Two Traditions

Language may serve a variety of functions, but it is generally agreed
that language is predominantly a tool for communication. Language
may be used to describe states of affairs and to express thoughts and
emotions; it may be used to manipulate people, to convince, to beg,
to demand, and so on. All of these functions of language must as-
sume a communicative context and must be constrained by the rules
of communication. In order to be understood, speakers and writers
implicitly observe such rules even when they "only" make assertions

or provide descriptions and even when no clear target audience is present. Listeners do their part by observing the same principles and assuming that speakers do so as well. This is a sketchy description of the cooperation that occurs between people when they are engaged in communication (Grice 1975). It does not necessarily tell us how the cognitive system operates in order to achieve such cooperation. One of the goals of cognitive science is to discover precisely those processes that underlie our use of language and to explain how they eventually afford communicative interaction.

Paradigms of linguistic research vary in terms of how central communication processes are to their theories. At one pole is the tradition in which communicative principles are assumed to be irrelevant to the study of language, such as formal theoretical linguistics, as exemplified by Noam Chomsky. In his *Russell Lectures*, Chomsky puts it succinctly: "Though consideration of intended effects avoids some problems, it seems to me that no matter how fully elaborated, it will at best provide an analysis of successful communication, but not of meaning or of the use of language, which need not involve communication or even the attempt to communicate" (Chomsky 1971: 19). This basic assumption has been part of all standard grammatical theories influenced by Chomsky's research program, regardless of whether their main concern is with the syntactic or the semantic aspects of language (cf. Lyons 1977).

At the other end of the continuum, standard pragmatic theories assume that in order to understand the way language works one must consider communicative principles, as exemplified by H. P. Grice's influential work on meaning and on the rules of conversation. Grice (1975) argued that the very act of using language raises expectations of cooperation between speakers and listeners. Listeners expect speakers to be clear, truthful, informative, and relevant. These expectations guide every process of interpretation, regardless of whether the speaker fulfills them or appears to be violating them. When Joel asks his roommate, Jim, what the date is and Jim answers,

Today is September 3d,   (1)

he is probably following the rules of conversation by being truthful, relevant, and so on. But if Jim says the same thing (1) without being asked about the date, he is flouting the rules of conversation. Joel may then infer that Jim wants to remind him that their rent check is overdue. In this way the very same expectations are exploited to convey the implied meaning of the utterance (Sperber and Wilson 1986: 37).

Even though the standard grammatical and pragmatic approaches are diametrically opposed on most issues, they share two basic assumptions about non-literal language. First, both assume the priority of the

literal. When, for example, a metaphorical sentence is encountered in discourse, its literal interpretation is always apprehended before any figurative interpretations are considered. Second, they both assume that a metaphorical interpretation is sought *only after* some anomaly, semantic deviance, or pragmatic rule violation has occurred. For standard semantic theories, semantic anomaly triggers a search for non-literal meanings. For standard pragmatic theories, apparent violations of discourse rules or rules of conversation trigger the search for non-literal meanings.

The assumption of the priority of the literal, then, is central to both types of theories. Standard semantic theories, which ignore communication, take it for granted that a literal interpretation is always the first step in the comprehension process. This assumption simplifies many aspects of those theories, but it complicates their accounts of metaphorical interpretation. For example, to explain how people understand such sentences as

a fisherman is a spider   (2)

these theories may need to postulate special, non-literal processing mechanisms. One way to resolve the problem is simply to avoid it. Donald Davidson (1979 [1978]), for example, argues that a semantic account of this sentence is simply that the literal meaning is false because a fisherman is not a spider. He suggests that any "meaning" beyond this, such as a possibly metaphorical one, is irrelevant to semantic analysis. A more traditional solution to the problem, however, does go beyond the initial literal analysis. Samuel Levin (1977) provides an elaborate theory of the semantic analysis of metaphor, again assuming that the initial interpretation of the sentence is always literal. A literal analysis initially yields semantic deviance for such sentences as (2) because several selection restrictions are violated. For example, fisherman has the feature "human," whereas spider does not. Such violations, according to Levin, are the crucial triggers whose activation leads to solving the comprehension problem. The violations trigger rules that adjust selection restrictions by changing certain features, resolving the anomaly, and ultimately yielding a metaphorical meaning. In this way, solely by semantic means, such as deviance and construing rules, and without any reference to aspects of communication, does this approach propose to account for metaphorical meaning.

A solution such as Levin's has its limitations. One drawback results from the assumption that deviance is a necessary condition for the construction of a metaphorical meaning. Thus, it cannot account for metaphors that do not violate semantic rules. For example, the sentence

it is a desert   (3)

does not involve any semantic anomaly or deviance, yet it can be understood metaphorically, as in *The major problem for many people these days is loneliness. It is a desert.* Such a theory as Levin's may be able to explain this in the following way: Once the indexical "it" is replaced with "loneliness," the result, *loneliness is a desert*, manifests the same kind of anomaly as does (2) *a fisherman is a spider*. This anomaly would then trigger the adjustment rules. Even though such counterexamples as (3) can be explained away by such a semantic theory, there are other kinds of sentences that are understood metaphorically, yet involve no semantic deviance. One example of this type is

Groucho's tongue is not a bayonet,   (4)

which is neither semantically deviant nor anomalous. In spite of lacking a "trigger" of deviance, this sentence does have a metaphorical meaning, which is so clear that most people will be quick to disagree with it. The semantic approach fails to account for such examples precisely because it ignores the communicative aspects of utterances.

Standard pragmatic theories, such as Grice's, are better equipped to deal with these examples, yet they make the same two basic assumptions. The first assumption is embodied in the distinction between sentence meaning and speaker meaning. Sentence meaning is supposed to be the basis for determining what the speaker means. As long as the speaker does not flout the basic "maxims" of cooperation, the speaker probably means what she says. In such cases, the sentence meaning is computed prior to the speaker meaning and is taken as what the speaker intended when there is no reason to believe otherwise. John Searle (1979) expresses the same idea when he suggests that the understanding system initially assesses how plausible a literal interpretation of each sentence would be. So, according to this notion,

the wind is a rocking chair   (5)

would first be analyzed literally, only to yield an anomaly. Given this anomaly, the next step would be to look for an alternative metaphorical meaning in the sentence. The anomaly, or deviance, is a necessary triggering condition for the construction of a metaphorical interpretation. This idea was clearly stated by Bess Altwerger and Steven L. Strauss (1987), who said that the "[communicative] exchange is liable to break down to the extent that anomaly is introduced," and, because "language users simply do not accept anomaly," an inference is required to "turn the deviance into meaningfulness." In this way, "metaphors may be viewed as arising from the apparent violation of one of Grice's conversational maxims" (ibid.: 180–81).

If the violation of Grice's maxims triggers such metaphorical interpretation, this approach can also account for sentences, such as (4) *Groucho's tongue is not a bayonet*, that yield a perfectly acceptable

literal interpretation, but are nonetheless understood metaphorically. In this case, the sentence is first analyzed literally, only to suggest the obvious: Groucho's tongue is excluded from the category of bayonets. This trivial interpretation violates the maxim of informativeness. Assuming that the speaker is being cooperative, one looks for an alternative, more informative, metaphorical interpretation. In this way, rule violation is a necessary condition for metaphor comprehension.

In summary, both standard semantic and pragmatic theories assume that the construction of literal meaning is always the first step in interpreting any sentence in any context. Furthermore, both approaches assume that metaphorical interpretations depend on specific triggering conditions. To be sure, they differ on the nature of these conditions, but they agree that a search for metaphorical meaning starts only when a literal interpretation is either anomalous or deviant, or when it violates one or another rule of conversation.

Communication and Metaphor: An Alternative

We will challenge the central assumptions of the standard theories and argue that they are inadequate for a psychological theory of metaphor understanding. Our central argument is that communicative principles are equally important for literal and metaphorical language use and that they constrain literal and metaphorical interpretations in the same way. We will discuss conditions under which literal interpretations seem to be constructed irrespective of communicative needs and will show that a metaphorical interpretation may be constructed in exactly the same fashion, thus suggesting that a "trigger" is not a necessary condition for understanding metaphors. Then we will consider those linguistic operations that seem to be constrained by communicative functions. For example, we will discuss the cases of apparent rule violation that lead to a metaphorical interpretation, arguing that the same conditions may also underlie the construction of literal interpretations. In this sense, both literal and metaphorical interpretations may be similarly constrained by communicative rules. After discussing the conditions for the construction of either literal or metaphorical interpretations, we will focus on *how* metaphors are understood. We will argue that communicative principles play a major role at this level of analysis and that they operate similarly for literal and metaphorical statements. We will then suggest that metaphorical statements may be distinguished from literal ones on the basis of what they are *used* for, instead of the conditions for their comprehension.

The Priority Issue

Competent speakers of a language cannot ignore the meaning of a word. People cannot, for example, name the color of the ink in which

the letters are written without also realizing the meaning of the word those letters comprise (Stroop 1935). Similarly, we cannot prevent ourselves from understanding the meanings of sentences. Semantic and syntactic operations simply take place without our conscious control, in an involuntary manner. Some aspects of understanding may take place, then, regardless of any communicative needs or circumstances. The question is, what is "the meaning" that we can't help but grasp? Both approaches that we presented above implicitly assume that the literal meanings of sentences are always obligatorily computed. Are metaphorical meanings grasped only optionally, then, or are they also computed obligatorily?

In order to address this question, Sam Glucksberg, Patricia Gildea, and Howard Bookin (1982) asked people to read sentences and to focus on their literal meanings. Subjects were asked to assess the literal truth of each sentence as quickly and accurately as they could. Some of these sentences were literally false, but made sense metaphorically. For example,

a mountain road is a snake   (6)

is literally false but metaphorically sensible. Such sentences were found to take longer to reject as false than literally false sentences that do not make sense metaphorically, such as

a mountain road is a tree.   (7)

A reasonable interpretation of such findings is that a natural agreement with the metaphorical meaning of a sentence such as (6) interferes with the task's demand for a negative answer. When such sentences had a metaphorical meaning, people could not help but grasp it. It seems, then, that one of the assumptions made by both approaches is unsupported. Literal interpretations need not take priority over alternative, metaphorical meanings. Metaphorical meaning seems to be as automatically grasped as the literal. In this sense, the construction of metaphorical interpretation may be as free of communicative constraints as is literal-language understanding.

Yet both standard approaches can defend themselves against such an interpretation of these empirical findings. But in order to do so, they must invoke the second assumption, stating that rule violation or semantic deviance is a necessary and sufficient condition for the construction of a metaphorical interpretation. They would argue, for example, as Marcelo Dascal (1987, 1989) did, that the literal interpretation of (6) was computed first and that this interpretation violated some rule. A semantic theory, such as Levin's, would suggest that a discrepancy between the selection restrictions of "mountain road" and "snake" results in deviance. A Gricean approach would focus on the

violation of the truth maxim. Regardless of the theory one chooses, one would have to agree that some rule violation or deviance could have been detected and that this violation could have been sufficient to trigger the construction of a metaphorical interpretation. If all this happens fast enough, the metaphorical meaning could interfere with people's responses to the literal meaning. So, one could explain the empirical results by assuming that the literal interpretation is indeed computed first, only to result in an anomaly, which, in turn, triggers a metaphorical interpretation.

Such an alternative explanation is reasonable and is perfectly consistent with both the semantic and the pragmatic approach. However, we believe that it is false for the following reasons: The main thrust of the alternative explanation rests on the second assumption, that is, the assumption of anomaly, which holds only because the sentences used in the experiment were literally false. But one can easily imagine an analogous pattern of results with sentences that are not semantically deviant. For example, one could ask people to assess the literal truth of such a sentence as

a mountain road is not a snake.   (8)

In contrast to (6) *a mountain road is a snake*, (8) obeys semantic rules. According to either of the standard views, then, such a sentence would not trigger the construction of a metaphorical interpretation. We would predict the reverse: the metaphorical meaning, again, cannot be ignored and will interfere with the ability to respond to the literal meaning. Yet this hypothetical experiment has two potential problems that could be used as counterarguments by standard semantic and pragmatic approaches. One obvious argument for the semantic-like approach is that since semantic deviance *is* a necessary condition for constructing a metaphorical meaning, when this condition is not fulfilled, as in the case of (8), pragmatic processes that take place further down the line may (or may not) eventually recover the metaphorical meaning. This would imply that we would *not* get the predicted results, that is, the potential metaphorical meaning would not interfere with grasping the literal. A Gricean response will also reject this experiment, but for a different reason. Even if it yields the predicted result, a Gricean principle can still account for it: Just as (6) violates Grice's truth maxim, (8) violates the maxim that utterances be informative. Because *a mountain road is not a snake* states the obvious, it may automatically trigger the construction of a more informative, metaphorical interpretation. Again, it seems that a metaphorical meaning must be triggered by a rule violation or semantic deviance. However, such a conclusion may be premature. We will present evidence to refute the objections of both standard semantic and standard pragmatic approaches.

Is Anomaly Necessary?

Up to this point, we have tried to argue against the assumption that literal meanings are computed prior to any other, non-literal meanings. We suggested that metaphorical meaning might be computed just as immediately and obligatorily as literal meaning. But we ran into a problem posed by the second assumption, the assumption that whenever deviance is detected in the literal interpretation, it triggers a metaphorical interpretation. The only way we can show that metaphors are just as hard to ignore is to show that they are computed even when no conceivable "trigger" exists. To demonstrate this, Boaz Keysar (1989a) presented people with sentences that could have both literal and metaphorical meanings. For instance,

Rena lives in a castle   (9)

may be taken literally, if she lives in an actual chateau, or metaphorically, if her home is an ordinary residence, but is also private, protected, and so on. Such sentences as (9) were read in contexts that provided information about the literal or metaphorical truth of each sentence. For example, one context described Rena as a retired rock star who bought a chateau in France, thus rendering (9) literally true; another described her as having bought a small adobe box of a house in the middle of the Arizona desert, thus making the sentence literally false. In addition to these details, the story context provided information about how private and secluded the place was. In one case, Rena was said to have complete privacy and luxury, never being bothered by groupies, which rendered the sentence metaphorically true. Alternatively, the information provided suggested that the sentence was metaphorically false by describing Rena's lack of privacy, due to a housing development's being built nearby, and so forth. In such contexts, then, a sentence like (9) had both a literal and a metaphorical meaning. In two of the cases, the two meanings were consistent with each other, when both were either true or false. But in two other cases, the two meanings contradicted each other, one being true and the other false. Again, people were simply asked to assess the *literal* truth of such sentences as (9), given the story context.

The question of interest for our current purposes is: Are people able to focus completely on the literal meaning when the sentence makes sense, or are they unable to ignore the alternative metaphorical meaning even where no triggering conditions exist? The results clearly showed that even when a sentence was literally true, people computed the metaphorical meaning. They were also slower to verify a literally true sentence when its metaphorical meaning was false. A second experiment (Keysar 1989a) used a different technique to show that these findings could also be extrapolated to normal reading practices (as opposed to the verification of sentences). We suggest, then,

that the construction of metaphorical interpretations need not require any triggering conditions at all.

We are now in a position to reject the two basic assumptions shared by the standard semantic and pragmatic approaches. First, literal interpretations do not take priority over metaphorical interpretations. Metaphorical meaning may be grasped as immediately and automatically as literal meanings. Secondly, we showed that neither a semantic anomaly nor the violation of a conversational rule is necessary for a metaphorical interpretation to be made. Recall that these two assumptions were the basis on which both standard approaches distinguished between literal and metaphorical interpretations: the metaphorical meaning, but not the literal, was considered the result of recognizing a communicative intent, that is, speakers must mean something else if anomaly is detected. However, what we have shown here suggests that the identification of both literal and metaphorical meanings may take place regardless of any communicative considerations.

When Are Communicative Principles Followed?

When a certain meaning makes sense in context, it is grasped whether it has a communicative function or not. Yet although a communicative function may not determine the detection of a sentence's potential meaning, it can guide the determination of which meaning is the intended one. In this sense, Grice's maxims may function as decision rules, if not as triggers for the construction of an interpretation. For example, a sentence may be taken metaphorically because of a judgment that the literal meaning could not have been intended. In the 1988 televised debate between the two U.S. vice-presidential candidates, Dan Quayle tried to assuage potential concerns about his age by reminding the audience that President Kennedy had also been young when he was elected. When his opponent, Senator Bentsen, responded, "You are no Jack Kennedy," his statement was taken figuratively because it was unreasonable to assume that he had meant it literally. In this sense, the determination that the metaphorical meaning is the more appropriate interpretation is guided by communicative principles. Is this a unique feature of non-literal speech? We suggest that literal interpretations are constrained by exactly the same kinds of communicative considerations. We argue that in analogous situations a statement may be taken literally because it couldn't have been meant metaphorically. In this sense, the decision rule operates similarly for literal and for metaphorical language.

To demonstrate this process, Keysar (1989b) asked people to complete and interpret such sentences as

if this place were not a prison, then——,   (10)

which presupposes that "this place" *is* a prison. These sentences appeared in different contexts, some of which were ambiguous and afforded both literal and metaphorical interpretations. Others included information to suggest that either a literal or a metaphorical reading of the presupposition would be inappropriate. For example, one context described the speaker as talking about a workplace, thus rendering the assumption literally false. Alternatively, another context provided information about how free and liberal the atmosphere was in "that place," thus suggesting that the presupposition was metaphorically false. The results were symmetrical for literal and metaphorical interpretations. In keeping with Grice's model, when the context rendered a literal interpretation of the presupposition of a sentence such as (10) false, people were much more likely to take the sentence metaphorically. More interestingly, when the context rendered it *metaphorically* false, people were more likely to take the sentence *literally*. The time needed to understand these sentences in their different contexts added convergent support for our hypothesis. People needed more time to understand a sentence metaphorically when doing so was the result of rejecting the alternative literal interpretation than when the metaphorical meaning was simply suggested by the context. Similarly, comprehension took longer for sentences that were interpreted *literally* because they could not have been meant metaphorically than for sentences whose interpretation did not involve such a decision. We conclude from these results that literal and metaphorical interpretations follow exactly the same kinds of decision rules when intended meanings are detected.

This conclusion implies a more general one: When communicative principles *are* taken into account in comprehension, they operate similarly for literal and metaphorical interpretations. When people understand language, they are constantly faced with the need to disambiguate utterances in order to arrive at intended meanings. Because literal language does not take priority in comprehension, this interpretation problem must be the same regardless of the final product of the interpretation, that is, regardless of whether the utterance was meant to be taken literally or figuratively. To solve this interpretation problem, people must use the same kind of strategy with both literal and metaphorical language.

How Is a Metaphorical Interpretation Constructed?

Until now our discussion has focused on metaphor identification. We have considered the conditions for metaphorical and literal interpretations, and the reasons for deciding which one is intended, but we have not yet directly addressed the problem of metaphor comprehension: *How* are metaphors understood? In what way is their inter-

pretation constructed? The theories discussed so far do not provide adequate answers to these questions. Some semantic approaches, such as Levin's, do outline the way that a metaphor may be interpreted as an extension of a literal interpretation via construing rules. Unfortunately, such approaches depend on the unfounded assumptions of the literal interpretation's primacy and the necessity of deviance for metaphor comprehension. Gricean theories do not address this question at all, but instead make the vague assumption that the listener or reader "searches" for an alternative metaphorical meaning once rule violation has been detected. Searle (1979) does provide an account of how metaphors may be constructed. He outlines a number of paths that may be taken from the literal to the metaphorical meaning in the process of comprehending a sentence. But, as with Levin's model, the empirical evidence refutes the basic assumptions underlying such a model of stages. Neither model, then, can serve as a basis for a psychological theory of metaphor comprehension. In the remainder of this paper we will point out the shortcomings of traditional solutions to this problem and then will outline an alternative model that accounts for the way metaphors are understood. We will show how communicative principles, such as relevance and informativeness, underlie the comprehension process, arguing that any adequate theory of metaphor understanding must accommodate those principles.

Metaphor as Comparison or Interaction

Consider the statement

inflation is a time bomb.   (11)

What does it mean to comprehend such a metaphor? A major part of understanding entails uncovering the implied relationship between the two terms, to answer the question: What relationship between "inflation" and "time bomb" is expressed in this sentence? This problem is easily resolved in the course of understanding such sentences as

a grenade is a time bomb.   (12)

Here, "grenade" is understood as a member of the category "time bombs" because the statement's form, "an A(NP) is a B(NP)," typically places "A" in the category "B." This poses a problem for interpreting (11), however, because the category "time bombs" does not typically include "inflation" as a member. A solution to this problem was offered by the classical theory of metaphor, Aristotle's comparison view (Roberts 1952). According to this view, such metaphors as (11) do not express a category relationship between "inflation" and "time bomb" but, instead, a comparison. Although this view was strongly criticized by Max Black in 1962, it still appeals to many. George Miller

(1979), for example, defended the comparison view and developed an interesting version of it. He argued that a "metaphor is an abbreviated simile" and that, because a simile is a comparison statement, "making the comparison explicit is the first riddle that a reader must solve" on encountering a metaphor. He suggests, for example, that understanding the metaphor

his wife is his mother   (13)

"requires the reinsertion of 'is like,' and the comparison so reconstructed provides the basis for understanding the sentence as a metaphor" (ibid.: 202). The basic idea of this approach is that when people understand metaphors they try to uncover the similarities between the two terms, the topic and the vehicle. Similarity, then, is the relation that underlies metaphors, so the metaphor *inflation is a time bomb* is understood as a metaphorical similarity statement, or simile:

inflation is *like* a time bomb.   (14)

Such similes were further analyzed by Andrew Ortony (1979) when he extended Amos Tversky's (1977) model of literal similarity. Ortony characterized the similarity relation expressed in a simile by contrast with such literal similarity statements as

a mine is like a time bomb.   (15)

According to this theory, comparing two such things involves identifying the sets of their common and distinctive features. Both mines and inflation, for example, share with time bombs such features as "may cause harm," "may have an effect at an unpredictable time in the future," and so on. Ortony argued that such literal comparisons differ from metaphorical comparisons in one crucial respect: the relative salience of the shared features. He suggested that in literal similarity statements the shared features are highly salient for both terms, while metaphorical comparisons always involve salience imbalance. In metaphorical comparisons, the shared features are highly salient for the vehicle, but have low salience for the topic. For example, the feature "may have an effect at an unpredictable time in the future" is highly salient for "time bomb" but not for "inflation." With this analysis of the salience of features, Ortony was able to account for anomalies as well. A statement such as

a book is like a time bomb   (16)

yields a set of common features (e.g., inanimate object) that is low in salience for both book and time bomb. As a result, comparisons of this sort are uninformative and thus are perceived as pragmatically anomalous. The same holds true for such reversed similes as

a time bomb is like inflation.   (17)

According to the salience-imbalance theory, these statements involve common features that are highly salient for the first term, "time bomb," but low in salience for the second term. Because they are uninformatively construed, these similes become anomalous once they are reversed.

While we agree with Ortony's appeal to communicative principles, such as informativeness, we argue that this is precisely where the salience-imbalance model falls short (Glucksberg and Keysar 1990). We need to distinguish property salience in the mind of the speaker from property salience in the mind of a potential hearer. A speaker who tells someone that

a nectarine is like a peach   (18)

may have properties in mind of high salience for both topic and vehicle: both nectarines and peaches are red-gold, juicy, taste alike, have pits, and so on. For a hearer who is already aware of these similarities, the statement is uninformative, so it would probably be taken non-literally because uttering such a statement violates the conversational maxim to be informative. For hearers who do not know what nectarines are like, however, the statement would be informative. At the same time, the statement involves properties of high salience for the vehicle, peaches, but of either low or no salience *in the mind of the hearer* for the topic, nectarine. Indeed, for hearers, all informative statements have this character: the ground must, by definition, be highly salient in the sentence vehicle and of either low salience or none at all in the sentence topic. Therefore, sentence (18) illustrates that comparisons with low/high matches are not uniquely metaphorical. Informative literal statements have the same structure.

This analysis reveals a more central problem with comparison models in general. If a metaphor is conceived of as a comparison, then hearers must know enough about the topic to discover the grounds for comparison. In Ortony's terms, the set of (relevant) common features must not be empty. But people could easily understand such similes as

Ceausescu is a small-scale Hitler   (19)

even before knowing about the atrocities committed by the Ceausescu regime in Romania. Such sentences demonstrate that people readily understand a metaphor even when the common feature set does not initially include anything that is relevant to the meaning of that metaphor.

Examples such as (19) motivated Black (1962) to reject the comparison view and to suggest an alternative. He argued that what is uncovered when a metaphor is understood is not a comparison or a similarity relation, but a more complex interaction between the two

terms. The notion of "interaction" is sufficiently vague to allow for a variety of interpretations, but the underlying idea is that the topic of a metaphor is perceived "in terms of" the vehicle. Even though we are sympathetic to the general idea of the interaction view, it is notoriously difficult to evaluate because no theory suggests a detailed analysis of such an interaction. For example, this view is unclear about what relationship between the topic and the vehicle must be constructed in order to understand the metaphor. In the next section, we will outline a new theory of metaphor (Glucksberg and Keysar 1990) that avoids the problems of the comparison view and provides a straightforward mechanism for uncovering metaphorical relations. After we describe the theory, we will demonstrate how it explains seemingly unrelated metaphoric phenomena.

Metaphor as Categorization

Earlier we noted that the comparison view assumes the relationship expressed in a metaphor to be that of similitude rather than categorization because the form of such metaphors as (2) *a fisherman is a spider* seems to suggest a categorical falsity. But this assumption holds true only if "spider" can refer solely to "any numerous arachnids of the order Araneae, having eight legs, a body divided into cephalothorax and an abdomen, and several spinnerets that produce silk used to make nests, cocoons, or webs for trapping insects" (*American Heritage Dictionary*, 2d coll. ed., s.v. "spider"). But if the vehicle stands for a higher-level category that can include "fisherman," then categorization may be what underlies such metaphors. We suggest that this is precisely the way metaphors are understood (Glucksberg and Keysar 1990). How is this achieved? Consider the following metaphor:

    his sermon was a sleeping pill.   (20)

We suggest that this metaphor is understood by placing "his sermon" in a category of sleep-inducing things, which is named "sleeping pills" and which may also include such things as certain lectures, lullabies, warm milk, and so on. In this sense, the metaphor is understood as exactly what it seems, a class-inclusion statement. In order to support this claim, we need to explain how this category is constructed and how it can be named "sleeping pills."

The construction of the category may be understood by considering the general nature of classification. Lawrence Barsalou (1983) has shown that multiple classification is a central aspect of the human conceptual system because everything can belong to numerous ad hoc categories. "Sleeping pills" may belong to many different classes of things: they are a kind of pill, but they are also a product; they belong

to the category of things that may be bought in pharmacies; they are in the class of sleep-inducing things, and so on. Similarly, "sermon" can be classified in many different categories, some of which also include sleeping pills as a member. We argue that in order to understand such metaphors as (20), precisely these categories must be uncovered, that is, the categories to which both topic and vehicle belong.

Yet the topic and the vehicle may each belong to numerous ad hoc categories, some of which may be totally irrelevant to the meaning of the metaphor. For example, both "sermons" and "sleeping pills" can belong to such categories as "should not be used excessively," which would not reveal the meaning of the metaphor. The question is, what determines which of these categories is relevant? One possible answer may be found in another aspect of ad hoc categories that Barsalou discovered. He has shown that, as with common categories, such ad hoc categories have a graded structure, meaning that some members are more typical than others. We suggest that the relevant category is the one that the vehicle typifies. Sleeping pills do not typify such categories as "products that may be bought in a pharmacy," but of all the categories to which sleeping pills belong, they uniquely typify the category of sleep-inducing things.

If metaphors of the form "an A is a B" are understood as genuine category statements, then "A" should be a subordinate of the superordinate "B." The second term should be the name of the superordinate category. But ad hoc categories do not have conventional names. They are normally descriptions, such as "things to pack when traveling." In contrast to such ad hoc categories, we argue that the category that is relevant to the metaphor does have a name: it takes the vehicle's name. We suggest that *because* the vehicle typifies the ad hoc category, it can be used as its name. Thus "sleeping pills" not only typify the category of sleep-inducing things, but may also be used to name that superordinate category. If the vehicle does indeed stand for the category, then the metaphor must be understood as expressing a genuine category relation.

The notion that a typical member, such as "sleeping pills," may actually name a category to which it belongs might initially seem strange. Yet, it is actually a device that is common to many languages. American Sign Language (ASL) does not have lexicalized terms for many superordinate categories. In order to refer to a superordinate category, signers use the terms for typical members of that category. For example, the sign for "jewelry" is "RING, NECKLACE, BRACELET, ETC." (Newport and Bellugi 1978). Only recently have some categories, such as "furniture" and "fruit," been lexicalized in ASL, probably because of the influence of spoken language.

Not only is the name of a typical member borrowed to name the

superordinate, but these two uses are kept clearly distinct. This sharp distinction was illustrated in a 1987 newspaper interview that was conducted in Israel. During that time, John Demjanjuk was on trial, accused of being "Ivan the Terrible," one of the sadistic guards at the concentration camp of Treblinka. Even though the main issue of the trial was whether or not he was Ivan the Terrible, Demjanjuk's name came to typify the category of people who committed such crimes against humanity. An interview with an Israeli citizen showed how the name "Demjanjuk" was distinctively used, on the one hand, for the person whose innocence/identity was in question and, on the other hand, to name the whole category of such people as Ivan the Terrible:

> Israeli: If he is a Demjanjuk, then he should be condemned to death.
>
> Reporter: But he is Demjanjuk, his name is John Demjanjuk.
>
> Israeli: I know his name is Demjanjuk, but I don't know if he is A Demjanjuk. (Shinoff 1987: 48)

The same linguistic device that is used in such languages as ASL may be responsible for the way we understand metaphors. (For examples of the use of this device in other languages, see Glucksberg and Keysar 1990: 8–9.) Just as a typical member of a category can name the superordinate in those languages, so can the vehicle of a metaphor name the category it typifies. Therefore, to understand such metaphors as (20) *his sermon was a sleeping pill*, we must construct an ad hoc category to which both belong and which sleeping pills typify. The relation that the metaphor expresses, then, is that of categorization: "his sermon" is classified as a member of this newly created category named by the vehicle.

This account is able to explain seemingly unrelated metaphoric phenomena by means of the same underlying mechanism. First, we will consider the nonreversibility of metaphors. If metaphors were understood as comparisons, it would be possible to reverse a metaphor and still have a meaningful sentence. But if

> the mind is a sponge   (21)

is reversed, the result is an anomaly:

> a sponge is the mind.   (22)

The reason is that (21) is not understood as a comparison but as a categorization statement, and, as with all categorization statements, such as

> a shark is a fish,   (23)

when reversed they yield an anomaly. Some metaphors, however, are still meaningful even after they have been reversed. But the reversed

sentence never retains the meaning of the original metaphor; it has a different ground. For example, while the statement

an encyclopedia is a jungle   (24)

suggests that it is very hard to find one's way through an encyclopedia,

this jungle is an encyclopedia   (25)

might mean that it is quite possible to learn a variety of things from this particular jungle. Those apparently reversible metaphors actually support our general claim: (24) is reversible because its topic, "encyclopedia," happens to typify a certain category, while the meanings of (24) and (25) are different precisely because "jungle" and "encyclopedia" typify radically different categories.

Interestingly, the same holds for metaphorical statements that look like comparisons, namely, similes. These do not behave like literal comparisons but, instead, like metaphors. When literal comparisons are reversed, the perceived similarity is normally reduced (Tversky 1977). For example, the statement

a quince is like an apple   (26)

expresses a stronger degree of similarity than

an apple is like a quince.   (27)

As Tversky observed, a variant of a category is generally more similar to a typical member than the other way around. In this sense, similarity statements are asymmetrical. If the relation uncovered in a simile were that of comparison or similarity, then reversing the simile should also result in a reduction of perceived similarity. However, similes do not manifest asymmetry but, instead, are nonreversible. As with metaphors, when similes are reversed they yield an anomalous statement. For example, the simile

the mind is like a sponge   (28)

is meaningless when reversed:

a sponge is like the mind.   (29)

The reason, we suggest, is that even with a simile that looks like a comparison, the relation uncovered is not that of comparison but of categorization.

This leads us to an additional metaphoric phenomenon that we can now explicate. If similarity is the underlying relation in similes, it is hard to explain why only similes and not literal comparisons can be expressed as class-inclusion statements. How can the comparison of two such dissimilar entities as "sponge" and "mind" be expressed not only as the comparison (28), but also as the class-inclusion statement (21) *the mind is a sponge*? In contrast to similes, literal comparisons between "like" things cannot be expressed in a class-inclusion form. For

example, a literal comparison, such as (26) *a quince is like an apple*, cannot be paraphrased **a quince is an apple*. The reason that similes can be expressed as metaphors is that they express an implicit relation of categorization. "The mind is like a sponge," then, can be expressed as a class-inclusion statement because "the mind is a sponge." In contrast to Miller's (1979) suggestion that metaphors are abbreviated similes, we argue that similes are taken as implicit metaphors, that is, as implicit classifications. This view of metaphor-as-class-inclusion allows us to reconsider the issue of metaphor and communication, but before doing so, we will assess an alternative approach in which verbal metaphors are assumed to almost always rely on preexisting conceptual metaphors.

Attributive Class Inclusion or Conceptual Mapping?

Our class-inclusion view suggests that metaphors involve the construction of novel categories that are typified by metaphor vehicles and that attribute the categories' properties to metaphor topics. George Lakoff and Mark Johnson (1980a, 1980b) argue that metaphorical language virtually always relies on preexisting conceptual mappings. They suggest that the human conceptual system is "fundamentally metaphorical in character" (1980b: 195). Abstract concepts, such as the concept of IDEA, are understood in terms of such conceptual mappings as IDEAS ARE ORGANISMS. This type of conceptual root metaphor is said to account for the systematicity that seems to underlie the way in which people talk about ideas. The metaphors

his ideas finally bore fruit,	(30)
their idea died on the vine,	(31)
she has a fertile imagination	(32)

are all motivated by the metaphorical mapping IDEAS ARE PLANTS, which is itself an instantiation of the more general mapping IDEAS ARE ORGANISMS.

Raymond Gibbs (in press) argues further that conceptual mappings form the basis for the vast majority of metaphors and that even apparently novel metaphors do not create attributive categories de novo, but instead rely upon preexisting conceptual mappings. The following metaphors, for example, presumably rely on the root metaphor LOVE IS A JOURNEY:

Our love is a voyage to the bottom of the sea.	(33)
We may have to go our separate ways.	(34)
Our love is a bumpy roller coaster ride.	(35)

When people use such metaphorical expressions, no new attributive categories are created. Instead, all are predicated on the same root metaphor of love as a journey: "There is a tight mapping according

to which entities in the domain of love (e.g., the lovers, their common goals, the love relationship, etc.) correspond systematically to entities in the domain of a journey" (ibid.). A closer analysis, however, reveals that this apparent systematicity is illusory. Consider the "journey" metaphor in the assertion that

our trip to Europe was a bumpy roller coaster ride.   (36)

If such metaphors derive their meaning from the preexisting conceptual mapping X IS A JOURNEY, then the meaning of (36) is simply that our trip to Europe was a journey! Surely, people would not use or interpret (36) in such an inane manner. Instead, the metaphor's vehicle, "bumpy roller coaster ride," is used to attribute such properties as "rough" and "uncomfortable" to the metaphor's topic, "our trip to Europe." Similarly, the roller coaster vehicle in (35) attributes to "our love" not the properties of a journey, but instead the properties of roughness, widely swinging moods and emotions, jolting ups and downs, and so on.

One might argue that a metaphor's topic interacts with its vehicle to produce a particular instantiation of a preexisting conceptual metaphor, but this is neither necessary nor desirable. The concept of "journey," for example, does not seem to be a factor in the interpretation of "bumpy roller coaster ride." Metaphors (35) and (36), above, seem to be similar to one another because of the specific properties that are typified by roller coaster rides, but not by journeys at all. In addition, (35) is more similar to (36) than it is to (33) or (34), above, even though all the love metaphors use one or another form of the LOVE IS A JOURNEY concept.

The preexisting-conceptual-mapping approach encounters severe difficulties in accounting for such relationships among metaphors precisely because it ignores the specific properties that metaphor vehicles attribute to their topics. We do not claim, of course, that such conceptual mappings as LOVE IS A JOURNEY do not exist, only that such general concepts may be irrelevant to specific metaphorical uses. In the extreme case, the conceptual-mapping approach may be downright misleading. We have already argued that the metaphor vehicle "bumpy roller coaster ride" may not depend on the "journey" concept at all. The generalization that follows from this example is that one cannot identify the ground for a metaphor from the common taxonomic category of the metaphor vehicle. Consider the following apparent instantiations of the LOVE IS A CONTAINER conceptual metaphor:

We're trapped in a rotten marriage.   (37)

Our relationship has become a prison.   (38)

Our relationship is nothing but a filing cabinet.   (39)

True, there are container-like properties in each of these metaphors: involuntary and unpleasant containment is obvious in (37) and (38), but what is one to make of (39)? When we asked some people to inter- pret (39), the overwhelming majority failed to mention any container- like properties at all, focusing instead on the organized, businesslike, unemotional properties associated with the metaphor vehicle, "filing cabinet."

The preexisting-conceptual-metaphor view, then, seems not only to be unable to account for how people interpret ordinary nominative metaphors, but it may also be seriously misleading. It cannot say how people interpret metaphors because it is silent on the specific proper- ties attributed to metaphor topics by specific metaphor vehicles. It may be misleading because it takes as the conceptual basis for a metaphor the common taxonomic category of the vehicle. As we have seen, this category may be totally irrelevant to the ground of a metaphor. What, then, are we to make of the availability of any number of metaphors of the form LOVE IS A [type of] JOURNEY? On the basis of our attribu- tive class-inclusion view, we would simply argue that various kinds of love can be identified with various types of journeys: pleasant, idyllic, stormy, placid, exciting, heavenly, fruitless, tormenting . . . the list is endless. To the extent that a type of journey, or a type of anything else, for that matter, exemplifies any one of these properties, could that type of journey be used as a metaphor for a certain type of love? Not if the concept of journey per se were totally irrelevant to the metaphor itself. Instead, certain aspects of love could be characterized as certain aspects of journeys or plants or animals, or anything else, to the extent that the "anything else" typified a property or properties that could be attributed to an aspect of love. At best, the preexisting-conceptual- mapping view provides a generalization: classes of metaphor topics, such as LOVE, may be characterized as one or another type of metaphor vehicle, such as JOURNEY. The more specific metaphoric work has to be done by more specific mechanisms, such as attributive class inclusion via the creation of novel metaphor/vehicle categories.

Communicative Functions of Metaphor

In our first sections, we considered the identification of metaphori- cal meaning and argued that some aspects of metaphor identification are governed by communicative principles, while some are indepen- dent of such principles. In contrast to standard theories, we suggested that literal language is subject to exactly the same constraints. We con- cluded that the processes underlying identification do not distinguish between literal and metaphorical language. We then went on to out- line a new model of metaphor understanding that does not rely on the assumptions of the standard theories. In this final section, we will con-

sider some functions of metaphor and propose that the class-inclusion approach has direct implications for a unique communicative *function* of metaphor.

Metaphors are considered tools for conveying new information by old means. But this function is not unique to metaphors. First, one can obviously convey new information by literal means. Secondly, people generally treat utterances as informative, and, as Grice has argued, even when certain utterances seem to be uninformative, they are eventually taken as informative. In addition, even when metaphors are no longer new, people still consider them to be metaphors. For example, statements as common as *my surgeon was a butcher* are still judged metaphorical. If the unique function of metaphor is not to convey new information, then what is its function? When do people choose to employ a metaphor rather than a simile or a literal form of expression?

We suggest that the unique function of metaphors does not lie in the novelty of the information conveyed, but in the way that it is conveyed. The information is conveyed via categorization, which suggests a strongly connected pattern of relations. This allows metaphors to function as descriptive devices that can capture the whole structure of a theme via a single phrase. For example, in a study of thematic similarity, a story about a businessman relying on advice from his bankrupt friend was entitled "The Blind Leading the Blind" (Seifert, Abelson, McKoon, and Ratcliff 1986). Such titles capture the essence of themes accurately and succinctly. Cristina Cacciari and Sam Glucksberg (1991) identified such phrases as quasi-metaphorical idioms and argued that they achieve their goal because they function in a manner that is analogous to a metaphor vehicle. Metaphorical idioms refer to particular events or actions and, at the same time, stand for the whole category of thematically related circumstances that they typify. In this example, the concrete actions of a blind person directing another and the failed businessman advising his unskilled friend are both members of a set of actions sharing this common theme that is best articulated by the metaphorical idiom *the blind leading the blind*. So by using a quasi-metaphorical idiom, one is able to distill a whole set of relevant aspects of a theme, an event, or an action. This is possible because the metaphor identifies the target event as an instance of a more general category of such events.

In a *communicative* setting, the choice of a metaphorical form indicates the strength of the suggested attribution. The more wholistic the intended attribution, the more appropriate is the metaphorical form. Consider the following example: In the film *The Great Dictator*, Charles Chaplin plays two characters, a dictator called Hynkel, who looks just like Hitler, and a Jewish barber who resembles Hitler (or Hynkel). The literal comparison statement

Hynkel was like Hitler   (40)

expresses certain similarities between Hynkel and Hitler. The comparison suggests, among other things, that Hynkel had eyes like Hitler's or the same type of mustache. The same dimensions of similarity underlie the literal comparison

the barber was like Hitler.   (41)

Both comparisons convey the information that each character resembled Hitler. In contrast, the metaphor

Hynkel was a Hitler   (42)

conveys a patterned set of attributes that characterizes what it means to be *a Hitler*. This patterned set of attributes is very different from the several independent dimensions of similarity that the literal comparison suggests. Even though both characters might have resembled Hitler, one could not truthfully say that

the barber was a Hitler.   *(43)

This difference between such literal comparisons and metaphors results from the class-inclusion nature of metaphors. According to the class-inclusion model, a metaphor is understood as the topic's classification in the category that is named by the vehicle. As a result of its classification, the topic "inherits" a whole set of features that indicates what it means to be a member of that category. The barber in *(43) cannot belong to that category of people typified by Hitler.

Now, compare the metaphor in (42) with the simile

Hynkel was like a Hitler.   (44)

The simile does not seem to have the same strength as the metaphor, which follows from the difference between similes and metaphors. In contrast to metaphors, similes involve *implicit* categorization. The "like" construction qualifies the class-inclusion nature of similes, and, as a result, the strength of the attribution may be qualified as well. Depending on whether speakers want to qualify their statements or not, they may choose one form over the other. For example, when one is talking about the character Hynkel, the metaphorical form (42) may be preferred over the simile (44). But if one is talking about the actor Chaplin, a qualified statement may be preferred:

when he played Hynkel, Chaplin was like a Hitler   (45)

may be preferred over

when he played Hynkel, Chaplin was a Hitler.   (46)

Some may disagree with our intuition about this and choose (46) over (45). But regardless of which form one decides to use, this choice will be a function of how strongly a categorization is to be conveyed.

Conclusions

Metaphorical and literal language are processed similarly. We have argued that communicative principles may not be required for the construction of a metaphorical interpretation any more than they are required for a literal interpretation. In this sense, the construction of metaphorical interpretations is no more a function of communication than is the construction of literal interpretations. Moreover, when communicative principles are applied in choosing an appropriate interpretation, they must operate similarly in the selection of either a metaphorical or a literal interpretation. In this sense, literal language is no less a function of communication than metaphors are. Communicative principles, then, cannot distinguish between the processes that lead one to take an utterance literally or metaphorically.

Although the processes that underlie comprehension do not distinguish between literal and metaphorical language, there is a difference in use, as we point out. We argue that a metaphor expresses a classification and is understood as such. As a classification, it suggests a strongly patterned inheritance of properties from the newly constructed attributive category to the topic. This distinguishes metaphors from both similes and literal comparisons. The use of a metaphor serves the communicative function of indicating the strength of the implied attribution. We suggest that, in general, the uniqueness of metaphors may be found in the realm of use, rather than in comprehension processes.

References

Altwerger, Bess, and Steven L. Strauss
    1987 "Readers' Understanding of Metaphor: Implications of Linguistic Theory for Psycholinguistic Research," in *Understanding Readers' Understanding*, edited by R. Tierney, P. Anders, and J. Mitchell, 175–88 (Hillsdale, NJ: Lawrence Erlbaum).
Barsalou, Lawrence
    1983 "Ad Hoc Categories," *Memory and Cognition* 11(3): 211–27.
Black, Max
    1962 "Metaphor," in *Models and Metaphors*, edited by M. Black, 25–47 (Ithaca, NY: Cornell University Press).
Cacciari, Cristina, and Sam Glucksberg
    1991 "Understanding Idiomatic Expressions: The Contribution of Word Meanings," in *Understanding Word and Sentence* (Advances in Psychology series), edited by G. B. Simpson, 217–40 (Amsterdam: North Holland Press).
Chomsky, Noam
    1971 *Problems of Knowledge and Freedom: The Russell Lectures* (New York: Random House).
Dascal, Marcelo
    1987 "Defending Literal Meaning," *Cognitive Science* 11: 259–81.
    1989 "On the Roles of Context and Literal Meaning in Understanding," *Cognitive Science* 13: 253–57.

Davidson, Donald
  1979 [1978] "What Metaphors Mean," in *On Metaphor*, edited by S. Sacks, 29–45 (Chicago: University of Chicago Press).
Gibbs, Raymond W., Jr.
  In press "Categorization and Metaphor Comprehension," *Psychological Review* 99.
Glucksberg, Sam, Patricia Gildea, and Howard Bookin
  1982 "On Understanding Nonliteral Speech: Can People Ignore Metaphors?" *Journal of Verbal Learning and Verbal Behavior* 21: 85–98.
Glucksberg, Sam, and Boaz Keysar
  1990 "Understanding Metaphorical Comparisons: Beyond Similarity," *Psychological Review* 97(1): 3–18.
Grice, H. P.
  1975 "Logic and Conversation," in *Syntax and Semantics*. Vol. 3, *Speech Acts*, edited by P. Cole and J. Morgan, 41–58 (New York: Academic Press).
Keysar, Boaz
  1989a "On the Functional Equivalence of Literal and Metaphorical Interpretations in Discourse," *Journal of Memory and Language* 28: 375–85.
  1989b "When Metaphor Fails: Contextual Control of Discourse Interpretation." Paper presented at the thirtieth annual meeting of the Psychonomic Society, Atlanta, GA.
Lakoff, George, and Mark Johnson
  1980a *Metaphors We Live By* (Chicago: University of Chicago Press).
  1980b "The Metaphorical Structure of the Human Conceptual System," *Cognitive Science* 4: 195–208.
Levin, Samuel
  1977 *The Semantics of Metaphor* (Baltimore: Johns Hopkins University Press).
Lyons, John
  1977 *Semantics* (New York: Cambridge University Press).
Miller, George
  1979 "Images and Models: Similes and Metaphors," in *Metaphor and Thought*, edited by A. Ortony, 202–50 (Cambridge: Cambridge University Press).
Newport, Elissa, and Ursula Bellugi
  1978 "Linguistic Expressions of Category Levels in a Visual-Gesture Language: A Flower Is a Flower Is a Flower," in *Cognition and Categorization*, edited by E. Rosch and B. B. Lloyd, 49–77 (Hillsdale, NJ: Lawrence Erlbaum).
Ortony, Andrew
  1979 "Beyond Literal Similarity," *Psychological Review* 86(3): 161–80.
Roberts, W. R., trans.
  1952 "Rhetoric," in *The Works of Aristotle*. Vol. 11, *Theoretica, De Rhetorica ad Alexandrum, Poetica*, edited by W. D. Ross (Oxford: Clarendon Press).
Searle, John
  1979 "Metaphor," in *Metaphor and Thought*, edited by A. Ortony, 92–123 (Cambridge: Cambridge University Press).
Seifert, Colleen, Robert Abelson, Gail McKoon, and Roger Ratcliff
  1986 "Memory Connections between Thematically Similar Episodes," *Journal of Experimental Psychology: Learning, Memory, and Cognition* 12: 220–31.
Shinoff, P.
  1987 "Demjanjuk War-Crimes Tribunal Strikes Deep Fear among Jews," *San Francisco Examiner* (14 June): A8.
Sperber, Dan, and Deirdre Wilson
  1986 *Relevance: Communication and Cognition* (Cambridge, MA: Harvard University Press).

Stroop, J. R.
   1935 "Studies of Interference in Serial Verbal Reactions," *Journal of Experimental Psychology* 18: 643–62.
Tversky, Amos
   1977 "Features of Similarity," *Psychological Review* 84(4): 327–52.