Beyond Literal Similarity

Andrew Ortony
Center for the Study of Reading
University of Illinois at Urbana–Champaign

Hitherto, theories of similarity have restricted themselves to judgments of what might be called literal similarity. A central thesis of this article is that a complete account of similarity needs also to be sensitive to nonliteralness, or metaphoricity, an aspect of similarity statements that is most evident in similes but that actually underlies metaphorical language in general. Theoretical arguments are advanced in support of the claim that metaphoricity can be represented in terms of the relative degrees of salience of matching (or matchable) attributes of the two terms in a comparison. A modification of Tversky's account of similarity is proposed. The implications of this proposal for similarity statements are discussed, along with implications for the psychological processes involved in their comprehension. It is argued that the general account of similarity proposed, including, as it does, nonliteral similarity, can form not only the basis of a theory of metaphor but can also give a credible account of the relationship between metaphor, analogy, and similarity.

Most theoretical approaches to similarity have been based on geometric models in which the similarity between two objects is a decreasing monotonic function of the distance between their representations in a multidimensional space. One problem with such approaches is the difficulty that they have in readily accounting for the lack of symmetry that is often found in similarity judgments, since geometric models are constrained by the fact that the distance between two points in a Euclidean space is the same regardless of the direction in which it is measured.[1]

Partly in response to this problem, Tversky (1977) proposed a contrast model, based on feature matching, that seems able to deal with the asymmetry problem. But neither Tversky's model as it stands, nor the spatial models that he criticized, are able to deal well with the radical asymmetry that is the

[1] There have, however, been proposals to avoid these difficulties, most recently by Krumhansl (1978).

hallmark of what might be called nonliteral similarity statements. The most obvious examples of such statements are similes, but nonliteral similarity seems to underlie many kinds of figurative uses of language and, in particular, seems to constitute the basis of metaphors.

Since it is widely believed that metaphorical language and metaphorical thought are based on statements and judgments of similarity (see Ortony, Reynolds, & Arter, 1978, for a review), there is a need for a theory of similarity that is sufficiently powerful to deal with nonliteral comparisons. The main purpose of this article is to offer some proposals that might lead to a comprehensive theory of similarity sufficiently powerful to constitute the basis of a theory of metaphor—in other words, a theory of similarity that goes beyond literal similarity to nonliteral similarity. Such a theory would also be tantamount to a theory of metaphoricity, since it would tell us what makes a metaphor metaphorical. What is presented is not yet an empirically tested theory. Several investigators are currently engaged in research that promises to provide data pertinent to some of the issues raised, but it may take many years to accumulate sufficient data to resolve them. Therefore, in the spirit of the hypothetico-deductive method, it seems appropriate to attempt to identify the issues and to propose possible solutions to them now.

Two theoretical constructs constitute central presuppositions of this article. Both are familiar, but it is important to lay them out so that their interpretation in the current context is clear. The first is that of a knowledge representation. The representations that will be presupposed here have been variously called *frames* (e.g., Minsky, 1975), *scripts* (e.g., Schank & Abelson, 1977), and *schemata* (e.g., Rumelhart & Ortony, 1977). Henceforth, the term *schema(ta)* will be used. A crucial characteristic of schemata is that they embed; a schema may contain tokens of, and tokens of it may be contained by, other schemata. Such subschemata can be viewed as representing predicates or attributes of the schemata that they dominate or by which they are dominated.

It is necessary that in any model dealing with the utilization of schemata in comprehension, the availability of schemata and subschemata should be sensitive to context. It is here that the second important concept, salience, comes into play. Several studies (e.g., Anderson & Ortony, 1975; Barclay, Bransford, Franks, McCarrell, & Nitsch, 1974) have shown that context influences and determines the particular aspects of word meanings that are salient on any particular occasion. In terms of schema theory, this means that in any particular context, some subschemata may be irrelevant or inappropriate, and consequently will not be involved in the comprehension process. This, in turn, can be expressed by saying that the salience of constituent structures in a knowledge representation can change as a function of context. In this article it will be assumed that salience can be operationally defined in terms of subjects' estimates of the prominence of a particular attribute with respect to a concept to which it does or could apply, although other (presumably highly correlated) measures, such as frequency or order of mention in an elicitation task, might do just as well.

The first major claim to be made is that Tversky's (1977) account of similarity can be profitably modified so as to provide a measure of similarity that is sensitive to metaphoricity. It will be argued that the principal source of metaphoricity is the difference in the relative salience of matching attributes of the terms in the comparison. This discussion can be found in the section on salience imbalance and metaphoricity. The next section identifies another variable that affects metaphoricity. This variable is attribute inequality; it arises when putative matches in fact only match metaphorically or at higher levels of abstraction in some assumed taxonomic structure. It is particularly noticeable when the domains from which the two terms are drawn are very disparate or incongruent. It is claimed that attribute inequality enhances the metaphorical effects of salience imbalance. The third section, on metaphoricity and asymmetry, discusses a number of symmetry-related issues arising from the proposed modification to Tversky's formulation of similarity. Then the next section, on two types of similarity statements,

attempts to show that similarity statements that are understood in terms of discovered matches of attributes have to be distinguished from those that can only be understood by introducing new attributes to the first term. In the last section, on metaphor, the implications of the proposals for a theory of metaphor are discussed.

Salience Imbalance as a Source of Metaphoricity

The theory of similarity proposed by Tversky (1977) is designed to account for the degree of judged similarity between two objects represented by, say, the terms *a* and *b*. The theory, which is well supported by the data, is that the perceived similarity, $s(a,b)$, is a weighted function of the intersection of attributes of *a* and *b* less the sum of a weighted function of the attributes distinctive to one and a weighted function of the attributes distinctive to the other, giving

$$s(a, b) = \theta f(A \cap B) - \alpha f(A - B)$$
$$-\beta f(B - A). \quad (1)$$

Here, the function f is a measure of the salience of features or sets of features. θ, α, and β are parameters that reflect the importance of the shared and distinctive features. Thus, for example, if $\alpha = \beta = 0$, and $\theta = 1$, the perceived similarity depends only on the shared attributes. Accordingly, changes in the values of α, β, and θ give rise to different similarity scales. *A* and *B* represent the sets of features of *a* and *b* respectively, while $(A - B)$ is the set of features distinctive to *A*, and $(B - A)$ is the set distinctive to *B*. It is assumed that the salience of a set of attributes is given by the sum of the salience of the members of the set. (See Tversky, 1977, p. 332.)

Tversky argued that two principal factors determine the salience of an attribute. The first, intensity, is independent of the object and

refers to factors that increase intensity or signal-to-noise ratio, such as the brightness of a light, the loudness of a tone, the saturation of a color, the frequency of an item, the clarity of a picture or the vividness of an image. (p. 342)

The second factor, diagnosticity, does depend on the object. Diagnosticity is concerned with the discriminability of an object from other objects with which it is implicitly or explicitly classified. It therefore presupposes a context of alternatives for the object. In the absence of such a context, or in a context in which the contrast set can only be considered to be the universe of objects in general, diagnosticity presumably plays no role. Unfortunately, Tversky did not explain how intensity and diagnosticity interact, but for present purposes the important point is that a consequence of Tversky's position is that when diagnosticity does not come into play, the salience of an attribute is independent of the object(s) of which it is an attribute. This means that the measure of an attribute's salience would be a constant and that it would contribute a constant amount to the overall salience of the stimulus. Mathematically this is convenient, since it means that the determination of the salience of the set of intersecting attributes in Equation 1 is quite straightforward: $f(A_n)$, the measure of attribute n in A, and $f(B_n)$, the measure of that same attribute in B, will be the same.

The chief difference between Tversky's model and the present one is that the latter requires that the salience of an attribute depends on the particular object of which it is an attribute, as well as on other, contextual, factors. In the present model, the perceived similarity between two objects depends, in part, on the relative level of salience of matching attributes; thus, in general, it cannot be assumed that $f(A_n) = f(B_n)$. Rather, it is supposed that independent of diagnosticity, an attribute can be more important with respect to one object than it is with respect to another, just as some members of natural categories are more typical than others (e.g., Rips, Shoben, & Smith, 1973; Rosch, 1975). For instance, as a rule, being made of iron is a more important attribute of magnets than it is of railroad tracks; and, even though not necessary for either, being red is a more important attribute of a fire truck than it is of a brick. The issue of whether or not the salience of an attribute is independent of the object that possesses it is an unsettled empirical question. Yet, it is not at all easy to test. It is difficult to distinguish between the absolute amount of sali-

ence an attribute contributes to an object and confounding variables such as the relative amount it contributes and the amount of knowledge that subjects have about the objects.

The rejection of Tversky's assumption that the salience of an attribute is independent of the object raises a difficulty for determining the salience of the intersection of A and B in Equation 1. How is $f(A \cap B)$ to be computed? The answer to this question constitutes one of the central claims of the present proposal: The salience of the intersection of A and B is dependent on the salience values of matching elements in B, rather than on some function of the values in both A and B or of their values in A alone. Equation 1 can now be rewritten as

$$s(a, b) = \theta f^B(A \cap B) - \alpha f^A(A - B) \\ - \beta f^B(B - A), \quad (2)$$

where f^A and f^B represent measures of salience based on the values in A and B respectively.

The measure of similarity, as given by Equation 2, remains essentially the same as Tversky's (Equation 1) in cases in which the matching attributes are of comparable salience in both A and B. So, when the matching attributes are of relatively low salience for both objects, those objects will not be judged very similar; when the salience levels of matching attributes are relatively high for both, the predicted similarities will again be similar to those from Equation 1. Henceforth, I shall refer to similarity statements in which the matching attributes are of comparably high salience (high B/high A) as *literal similarity statements,* although it should be recognized that this is a shorthand way of saying that such statements are predominantly literal, since literalness (and its complement, metaphoricity) is a matter of degree. This seems to accord with common sense. It could be taken as axiomatic that if two things share some characteristics that are important to both, then those things will be perceived as literally similar.

However, it is now possible to go further than the mere characterization of literal similarity. Both nonliteral similarity statements

and anomalous ones can be characterized. To the extent that matching attributes are of lower salience in A than they are in B, while there exist high salient attributes in B that cannot be applied to a, comparisons between the corresponding terms will be nonliteral (high B/low A). Such statements are usually called similes. And, to the extent that similarity statements are neither literal nor nonliteral in the sense just explained, they will be anomalous. Thus, literal and nonliteral similarity statements do not form mutually exclusive classes of statements. Nor, for that matter, do anomalous and meaningful ones. It is preferable to think in terms of three components of similarity statements: the literal, the nonliteral (or metaphorical), and the anomalous. Sometimes one or two of these components contribute virtually nothing to the perceived similarity (e.g., the anomalous component contributes nothing if a statement is perceived as a literal similarity statement). As will be discussed in greater detail with the help of Statements 1–4 below, the present proposal identifies two rather different sources of anomaly in similarity statements. One arises from the assertion of similarity between two things for which the only matches are of trivial attributes (low B/low A matches), and the other arises from cases in which there is, essentially, a reversed simile (low B/high A).

What is being claimed is that the imbalance, $I(a, b)$, in salience levels of matching attributes of the two terms is a principal source of metaphoricity. A convenient way of conceptualizing this imbalance is to visualize the attributes of a and b as a list with the most salient attributes at the top. Then salience imbalance can be thought of as the degree of diagonality from attributes in B to attributes in A and can be characterized, to a first approximation, by considering the combined effect of the difference in salience between the matching attributes for a and for b together with the (independent) degree of salience in each, as in Equation 3

$$I(a, b) = g[f^B(X_i, \ldots, X_{i+n}) \\ - f^A(X_i, \ldots, X_{i+n})], \quad (3)$$

where g is some, probably additive, function, and $\{X_i, \ldots, X_{i+n}\} = A \cap B$. In cases in

which salience imbalance is the only source of metaphoricity, the degree of metaphoricity can be equated with the degree of salience imbalance, or diagonality. It then follows from Equation 3 that the more literal a similarity statement is, the lower will be its degree of metaphoricity because the matching attributes have comparable (predominantly high) degrees of salience in both A and B. Furthermore, the terms in highly literal similarity statements are likely to be judged more similar than the terms in more metaphorical ones because the set of intersecting attributes is likely to be larger in the former; and the rated similarity of nonliteral similarity statements is likely to be higher than for anomalous ones because in the former the measure of salience of the set of intersecting attributes will be greater, as determined by the term $f^B(A \cap B)$, than that measure for anomalous comparisons, in which the intersecting set may often even be empty.

The present proposal, then, not only distinguishes literal from metaphorical similarity statements and provides a measure of similarity that is sensitive to both, it also characterizes two sources of anomaly in putative similarity statements. Consider the following similarity statements:

1. Billboards are like placards,

2. billboards are like warts,

3. billboards are like pears,

4. chairs are like syllogisms,

5. sleeping pills are like sermons.

According to the present proposal, Statement 1 is basically a literal comparison, since billboards and placards share a number of high-salient attributes. By contrast, Statement 2 is a nonliteral comparison because, although no high-salient attributes are shared, there are some high-salient attributes of warts that are less salient attributes of billboards (e.g., they are ugly), whereas there are other high-salient attributes of warts that cannot be applied to billboards at all (e.g., they are found on the skin). Thus, Statement 2 is a metaphorical similarity statement—that is, a simile. The remaining cases are anomalous. In Statement 3 the only attributes common

to both terms are trivial attributes low in salience, such as "being a thing" or "being a physical object," and thus the measure of similarity will be very low. What makes such a statement anomalous is that it cannot normally be expected to fulfill a communicative function. It seems to violate the Gricean maxim of quality (Grice, 1975). In Statement 4 the two terms seem not even to share low-salient attributes; consequently, it too is anomalous in essentially the same way. Finally, in Statement 5 the only match seems to be of an attribute like "being soporific," that is, a high-salient attribute of the a term and a low-salient attribute of the b term. The anomaly here is of a different kind. There is indeed a match, but both the similarity and the metaphoricity are low since the value of $f^B(A \cap B)$ depends on the (low) salience value of the attribute for the b terms, sermons.

One prediction that follows from this is that other things being equal, cases like Statement 5 will tend to receive slightly higher similarity ratings than cases like Statement 4, because the intersection term is necessarily always empty for the latter. Cases like Statement 5 might also be expected to receive higher similarity ratings than cases like Statement 3 because the abstract and low-salient attributes that are shared (e.g., "is a thing") by the terms in cases like Statement 3 are probably not directly represented in the schemata associated with the terms. By contrast, preliminary data that we have collected suggest that subjects often recognize that reversed similes, like Statement 5, are indeed reversed similes. Perhaps for this reason, rated similarities for reversed similes tend to be higher than predicted by Equation 2.

There is always the possibility that a reversed simile may give rise to a new simile with a totally different meaning but with comparably high metaphoricity and similarity ratings, as, for example, is the case with Statement 2 and its reversal, Statement 6:

6. Warts are like billboards.

Some people find that ugliness, although important in the meaning of Statement 2, be-

comes much less important in Statement 6, in which attributes related to prominence seem to come more into focus. Another example of the meaning change associated with the reversal of terms in a simile is provided by Statements 7 and 8:

7. Butchers are like surgeons,

8. surgeons are like butchers.

Here, not only are the bases of the comparisons different, they are actually incompatible. The present account explains this fact quite easily; it maintains that the basis of a similarity statement depends on the salience levels of matching attributes relative to the *b* term, and that due to salience imbalance, the basis is likely to be different as the two terms are interchanged. Notice, also, that in Statements 7 and 8 there is a reversal of affective valence. Affective attributes are often (although by no means always) an important part of the basis of nonliteral similarity statements.

In spite of the peculiar behavior of reversed similes, the similarity statements that the present proposals identify as anomalous are relatively difficult to interpret; the generation of an interpretation cannot normally be readily achieved. It is true, as with most well-formed sentences in a natural language, that an interpretation can be forced. It is usually possible to construct a context in which comparisons like Statements 3, 4, or 5, can be interpreted. It may be, however, that such a context only permits an interpretation by causing a change in the relative degrees of salience of the attributes of the terms. The point is not that is impossible to conjure up a suitable context—it almost never is impossible. The point is merely that it is much more difficult to produce such a context for anomalous cases than it is for meaningful ones.

What we have so far, then, is the proposal to replace Tversky's (1977) measure of salience of shared attributes with a measure based on the salience of those attributes in *B*. The contrast in salience levels of attributes in *B* and *A*, the diagonality, can then form the basis of a measure of metaphoricity

(Equation 3). This account seems to fulfill the need for a measure of similarity that does not predict unrealistically low similarity ratings for similes, as does Equation 1. However, salience imbalance is not the whole story.

Attribute Inequality in Incongruent Domains

The notion of matching attributes is obviously essential to the account of similarity that is being proposed. Both similarity and metaphoricity depend on it. Yet the notion is not without its problems, particularly if one tries to use it in the context of schema-theoretic representations. This section starts with a discussion of these problems and then examines their effects on metaphoricity.

Consider Statement 9:

9. Blood vessels are like aqueducts.

According to the present proposal, the basis of similarity lies in the matching attributes. At first glance it might seem that there is no problem; blood vessels and aqueducts are, at least, both channels through which liquids move. This, however, is a misleading oversimplification. It is true that aqueducts are channels for carrying liquids, but it is not very convincing to argue that some predicate like "is a channel for carrying liquids" is represented as an important part of a person's knowledge about aqueducts. It seems more natural to suppose that what is represented directly in the schema for an aqueduct is that it is a channel for carrying water. The predicate about liquids is the result of an inference through the one about water. Similarly, blood vessels convey blood; to be sure, "being a liquid" is an important attribute of blood, but presumably it does not appear as a subschema in the schema for blood vessels. But this seems to lead to the conclusion that Statement 9 is trivial, being based merely on the fact that both blood vessels and aqueducts are channels. The problem is that this conclusion is plainly wrong.

Suppose, for a moment, that instead of conceiving of attributes as simple predicates, we concentrate on schematic structure, taking into account the relationships between the

attributes, not just the attributes themselves. For *aqueduct,* it might be supposed that the schema is a structure resembling the following:

Structure 1

Aqueduct Schema—isa (x, Aqueduct)
    A1: isa (x, channel)
    A2: flows through (water, x)
    A3: purpose of (A1, A2)
    (et cetera)

This can be compared with Structure 2, which might be used to represent (some of) a person's knowledge about blood vessels:

Structure 2

Blood-Vessel Schema—isa (x, Blood Vessel)
    A1: isa (x, channel)
    A2: flows through (blood, x)
    A3: purpose of (A1, A2)
    A4: isa (artery, x)
    (et cetera)

If A1, A2, and A3 are viewed as attributes, then the match on A1 represents the fact that both are channels. But, in addition, although the variables in A2 differ (water in one case and blood in the other), the structure of A2 is the same in each case; in other words, the attributes, although not identical, are themselves similar. A recursive mechanism for finding matching attributes could thus reveal a second-order match of, for example, "flows through (liquid, x)." It is also important to notice that A3 is shared because A3 can be considered to be a kind of meta-attribute that incorporates information about interattribute relations (i.e., it incorporates information about the structure of the schemata themselves). This is one of the reasons why it is helpful to use an approach to knowledge representations that incorporates structure, an approach that is richer than a mere listing of simple attributes. It appears to be a good way of capturing the fact that many statements of similarity depend on some structural isomorphism between the knowledge associated with the two concepts rather than on merely a match of simple attributes. These observations indicate how sensitive attribute matching is to the way in which attributes are represented, how they relate, and what they are; they

may well help to establish the superiority of one kind of representational approach over another.

But we are not yet out of the woods. It can still be objected that the attributes that are purported to match do not really match at all. For example, in what sense are both blood vessels and aqueducts really channels? There are many kinds of channels, some man-made and varying greatly in such properties as size and material of construction, and some natural, also varying greatly in comparable properties. Furthermore, there is a wide variety of things that flow through channels—all sorts of liquids, gases, and even information. It is certainly the case that a blood vessel is a very different kind of channel from an aqueduct. But, if this is so, on what basis can it be claimed that "being a channel" is the *same* attribute for both? This is the attribute inequality problem. Attributes that may be nominally the same often seem to change their meanings as they are applied to terms in different domains, so that across those domains they are related by similarity rather than by identity.

The solution to the attribute inequality problem lies, at least partially, precisely in this last statement of it. The criterion for a match has to be that the attributes should be highly similar, rather than identical. This, of course, introduces a recursive element into our account of similarity, but there is nothing wrong with that. Returning to Statement 9, it can now be seen that Attributes A2 from Structures 1 and 2, although not identical, are highly similar, and that consequently they can figure in the set of intersecting attributes. Of course, the reason they are similar is that blood and water themselves share the high-salient attribute of being a liquid. So, one of the reasons that blood vessels are similar to aqueducts is that they share attributes of attributes.

There remains a difficult problem, however. The existence of higher order matches is not sufficient to ensure similarity. For example, penguins are not normally thought of as being similar to wolves, even though there is a higher order match—they are both animals. Why is that not sufficient to make

them similar? One reason is that it seems to be the only basis for a match, so that the distinctive attributes easily outweigh it.

Another reason, however, might be that there exist constraints on the level of specificity of the categories to which the things being compared are typically thought to belong. Although this proposal is very tentative, suppose it were assumed that in the absence of direct, first-order matches, the existence of some reasonably specific shared domain were a precondition for a sensible similarity statement; that is, that the existence of such a domain was a necessary but not a sufficient condition for two things to be perceived as being even potentially similar. One way to interpret the notion of a reasonably specific shared domain would be in terms of basic level categories (e.g., Rosch, Mervis, Grey, Johnson, & Boyes-Braem, 1976). A level of specificity at or below the basic level would count as being reasonably specific, whereas a level of specificity more abstract than the basic level would not. Or, one might appeal to a more sophisticated "natural" taxonomy wherein there are more levels from which to select a cutoff point, as with, for example, the taxonomy described by Berlin, Breedlove, and Raven (1973).

The consequences of this kind of assumption can now be considered. In an anomalous case like Statement 3, the domains of advertising and fruit do not come together in any conceivable taxonomic structure in a category at or below the specificity of a basic level category. So Statement 3 fails to satisfy the precondition for a sensible similarity statement. By contrast, it has been supposed that blood vessels and aqueducts come together in a domain of channels or conduits, which does satisfy the precondition. Penguins and wolves only meet in the animal category, which is at a level of specificity more abstract than the basic level; consequently, the precondition is again not satisfied. It need not necessarily be possible to find a natural, already encoded category; sometimes a psychologically plausible category has to be "constructed." This, for example, is probably the case for Statement 10:

10. Encyclopedias are like gold mines,

in which lexical descriptions of the intersecting category might be a "place where things are stored," a "place where things can be found," or a "source of utility." It is also one reason why, under suitable contextual conditions, superficially anomalous comparisons may become interpretable.

Whether or not the general problem of attribute inequality can be solved along these lines, it is a problem that has to be dealt with in any theory of similarity that relies on attribute matching. The question then arises as to how attribute inequality relates to metaphoricity. Does it, perhaps, enhance salience imbalance? It seems reasonable to suppose that it does. This would mean that less salience imbalance would be needed to give rise to some particular degree of metaphoricity if the matching attributes were not in fact identical. More obvious, however, is the possibility that metaphoricity is increased in this way if the matching attributes are related by nonliteral similarity. This situation is evident in Statement 10. When subjects are asked to list attributes of encyclopedias and gold mines, they often list "being valuable" as a high-salient attribute of both. Clearly, however, the sense of being valuable that applies to encyclopedias is one pertaining to intellectual or mental domains, whereas the sense that applies to gold mines is a financial or pecuniary one (cf. Schank & Abelson's, 1977, distinction between MTRANS and PTRANS). It might well be argued that the basic (literal) sense of "valuable" applies in the pecuniary domain and that the applications of the attribute in the intellectual domain is derivative or metaphorical.

To summarize, the domains from which the terms in a statement of comparison are drawn can be incongruent or semantically remote from one another. When this happens, there often appear to be no real attribute matches at all. To the extent that there are matches, they tend to depend not on attributes, or configurations of them, being *identical*, but on their being *similar*. This similarity itself often has an element of metaphoricity associated with it, as seems to be the case with "is valuable," and as is cer-

tainly the case with "involves digging around" in Statement 10. Thus, it seems that the notion of an attribute match has to be revised. It has to be based not on attribute identity but on attribute similarity. A match of attributes requires either high similarity between them or high metaphoricity as defined in Equation 2. It may well be that even when there is attribute inequality, there is still some degree of diagonality, but that perhaps domain incongruence enhances its effects on metaphoricity.

Domain incongruence and attribute inequality are of general importance in lexical semantics and especially in the analysis of dual function terms like *cold, bright, hard,* and so on. Consider Statement 11, discussed by Searle (in press):

11. Sally is (like) a block of ice.

One aspect of Statement 11 that is rather important, and sometimes overlooked, is its ambiguity. If it is used in the context of Sally coming in from an extremely cold environment, it will have a much lower degree of perceived metaphoricity than if it is used in the context of a disillusioned would-be lover bemoaning Sally's unresponsive frigidity. The reason for this difference is that in the first interpretation there is domain congruence, whereas in the second there is not. Thus, in the first interpretation, although being physically cold (e.g., to the touch) is not a high-salient, persevering attribute of Sally, it is a high-salient attribute of a block of ice; consequently, it satisfies the conditions for being a simile. The attribute *cold* applies to each term in the same domain. The communicative success of such a hyperbole depends on the fact that the intensity of the coldness is different vis-à-vis the two objects, but it does seem to be the same attribute.

By contrast, in the second interpretation, the attribute *cold* has to be applied across domains, namely from the physical, temperature domain appropriate to ice in the one case to the emotional domain applicable to "psychological" characteristics in the other. This should not be construed as a similarity statement that depends for its success on some kind of pun, for as has already been implied, this characteristic turns out to be quite widespread, often relying on systematic, conventional, underlying analogies between two different domains (in this case, temperature and emotions, and in others, luminosity and personality, etc.). In fact, Jackendoff (1975, 1976) claimed that the entire semantic system of English can be built up using such notions. He argued, for example, that *give* is basically the same verb as *go*, except that the domain of the former is possession, whereas that of the latter is location. Implicit in this approach is the view that nonliteral similarity is a fundamental building block of language in general.

It seems, then, that metaphoricity depends not only on an inequality between relative salience levels for cases in which there are genuine attribute matches but also on domain incongruence for cases in which there are not. One might imagine an experiment in which ambiguous comparisons like Statement 11 appeared in contexts that forced either the domain-congruent interpretation or the domain-incongruent interpretation. Subjects asked to rate the degree of metaphoricity and similarity should give higher metaphoricity ratings in the domain-incongruent interpretation but higher similarity ratings in the domain-congruent interpretations.

So, returning to Statement 10, when we say that encyclopedias are like gold mines, a cognitive gear change is needed. Gold mines are sources of physical wealth; encyclopedias, of mental wealth. The possibility of applying terms like *wealth* in two domains is a result of the fact that there are underlying attributes that are shared by the two applications of the term. The knowledge that the terms are conventionally so applied serves, in comprehension, to short-circuit what might otherwise have been a recursive process required to uncover those similarities. However, the domain incongruence serves to increase the semantic distance that the comprehension process has to bridge, so that the perceived metaphoricity of a similarity statement that involves fundamentally different domains will be greater.

Metaphoricity and the Asymmetry of Comparisons

In the present account of metaphoricity in similarity statements, a great deal of emphasis is placed on salience imbalance. This section deals with the relation of salience imbalance to a number of problems connected with the asymmetry of similarity statements. Factors resulting in changes in judged similarity and metaphoricity are also discussed. First, however, it will be helpful to review briefly an important component of Tversky's (1977) proposals, the diagnosticity principle. This principle is basically concerned with the fact that context can influence the salience of attributes. Indeed, the influence of context may even extend to introducing an attribute that otherwise would be trivial. Tversky's example is that the attribute *real* has no diagnostic value in the context of actual animals; that is, it would be of very low salience. Yet in a context of animals that included imaginary and mythical beasts, it might become very important.

If the proposed modification (Equation 2) to Tversky's model is accepted, it follows that a temporary (or, for that matter, a permanent) change in the salience levels of matching attributes should result in changes to judged metaphoricity and similarity. For example, many people have the intuition that specifying or linguistically constraining a dimension of similarity in what would otherwise be a very metaphorical similarity statement, reduces the degree of metaphoricity. Consider Statements 12 and 13:

12. John's face was like a beet,

13. John's face was red like a beet.

In Statement 13, John's face is compared to a beet with respect to redness. The effect of specifying the dimension is to identify, or "foreground," the most diagnostic attribute(s). Much the same would be true, although perhaps to a lesser extent, if *was like* in Statement 12 were changed to *looked like*. The consequence of foregrounding in Statement 13 is that all other attributes of both John's face and of beets have less impact on the perceived similarity between the two.

Another way of putting this is to say that the salience of the color attributes is increased above the salience of all the other attributes so that the latter no longer play a significant role. The result is a match of high-salient to high-salient attributes. Accordingly, judged metaphoricity should diminish from Statement 12 to Statement 13. Furthermore, the prediction that perceived similarity increases follows readily. However, even with respect to an individual attribute such as color or size, the match may not be perfect; John's face was perhaps not literally the color of a beet. This suggests that fine tuning is required, that the attribute of color itself has attributes that may be more or less well matched (e.g., intensity, hue, and saturation). Consequently, even when an attribute of comparison has been foregrounded in this way, the similarity of two objects with respect to that attribute can vary. In other words, attribute inequality can be found in literal similarity statements too, which is why it was suggested earlier that some degree of salience imbalance is probably a necessary condition for metaphoricity.

Even though the replacement in Statement 13 of a high B/low A salience match by a high B/high A match as a result of the explicit mention of a shared attribute has been proposed as an explanation of the reduction in perceived metaphoricity, this is not sufficient to guarantee symmetry. There remain strong constraints on the natural order of terms. One reason for this may be that whereas the attribute is matched qualitatively, it is not matched quantitatively. Perhaps statements like Statement 13 depend for their effectiveness on the intensity of the matching attribute being higher in B than in A. This may relate to Tversky's observation that we normally find the more natural order of terms in a similarity statement to be the one in which a deviant object is referred to in subject position, and the more prototypical one in the object position, as in Statement 14:

14. North Korea is like Red China.

In the case of Statement 13, then, something is needed for the b term that is more

prototypically red—it would be unnatural to compare the redness of John's face to something that was not generally believed to be typically red (e.g., a grape). This point becomes more obvious with attributes like cold, in which the perception and measurement of intensity is more commonplace.

The proposal, then, is that one reason why finding a nominal match, even if it is a high/high one, does not guarantee symmetry is that a matching attribute may vary with respect to some of its own attributes. Thus, in spite of the fact that "John is strong like an ox" isolates strength as the matching attribute, it is presumably the case that although both may be very strong, John is really very much less strong than the typical ox. In general, the determination of such within-attribute similarity calls for the same kind of operations that are required for normal similarity judgments, which again leads to the conclusion that the process of similarity perception may have to be viewed as a recursive one. As long as attributes are considered as subschemata, the idea of attributes themselves having attributes seems to be perfectly acceptable, since it is part of the theory of the representation of knowledge that the current account of similarity presupposes. However, it should be noted that there seems to be a price for gaining the flexibility that the embedding characteristic of schemata endows on attributes. It appears to be increasingly difficult to pin down the notion of an attribute.

In considering the question of symmetry, it is important not to overlook the fact that the sentence topic itself imposes constraints, so that generally a difference between "a is like b" and "b is like a" will always remain because in the first case the sentence topic is a and in the second it is b; that is, there are constraints resulting from such things as the given/new relationship and the subject/predicate relationship. Tversky refers to this as the focusing effect, noting that greater weight is normally assigned to the attributes of the subject term than to those of the second term (i.e., $\alpha > \beta$). If one considers the intersection term in Equation 2 as providing the basis of similarity, the question of symmetry seems to reduce to the fact that a

similarity relation can be symmetrical only if the basis for the comparison is the same, regardless of the order of mention of the terms in it. Thus, if a is like b in exactly the same respects in which b is like a, then the relation could be symmetrical. It should by now be clear that this condition can never hold for nonliteral similarity statements, because they share high-salient B and low-salient A attributes. If all the shared attributes in a simile have this high B/low A relation, then the simile, if reversed, will result in an anomalous comparison. Furthermore, it follows that with certain kinds of anomalous cases (low B/high A), the reversal will result in an interpretable simile, as was the case, for example, with Statement 5. Anomalous cases of low B/low A are also reversible in the trivial sense that they are uninterpretable in both directions for the same reasons.

The fact that similarity statements in general are not purely literal or purely nonliteral also has implications for symmetry. As mentioned earlier, the basis of Statement 2 lies in high-salient attributes of warts, such as being protrusions and being ugly. However, when reversed to give

15. warts are like billboards,

other attributes seem to take over; now the notions of prominence and obviousness seem to be more central. Thus the meaning of Statement 2 is different from that of Statement 15. In literal comparisons the basis of comparison is more likely to remain the same regardless of the order of the two terms, consequently the difference in meaning between the two orders is generally much less dramatic, although residual matches of high to low may still have an effect. This relates to Tversky's observation that "the variant is more similar to the prototype than the prototype is to the variant" (p. 333), as evidenced by comparing Statement 14 with its reversal. However, for cases in which (most) people perceive the matches as being at similar levels of salience, the maintenance of symmetry is easier, as in Statements 16 and 17:

16. Czechoslovakia is like Hungary,

17. Hungary is like Czechoslovakia.

The general conclusion, then, is that the degree of symmetry is inversely related to the degree of metaphoricity, so that the more metaphorical the comparison, the less symmetrical it is likely to be. Notwithstanding this, it remains true that other factors also cause asymmetry. For example, there may be subtle meaning changes resulting from high B/low A matches becoming low B/high A matches, and vice versa. Second, the kind of variables mentioned earlier—the subject/predicate relation, the given/new relation, relative amounts of knowledge associated with the terms, and the typicality of the terms—will almost always have a residual effect that can to a large extent be handled by accepting Tversky's account wherein, usually, $\alpha > \beta$, thereby giving less weight to the distinctive features of the second term. Finally, metaphoricity can be reduced when linguistic devices are used to foreground certain attributes or classes of them at the expense of others. This is why specifying a dimension with respect to which two otherwise substantially dissimilar objects are to be compared seems to reduce the metaphoricity. This is also why verbs of similitude more specific than *to be like,* that is, verbs such as *to look like* and *to sound like,* generally give rise to more literal similarity statements.

Two Types of Similarity Statement

The account of literal and nonliteral similarity statements that has been proposed locates the basis of comparison of two putatively similar terms in the intersection term of Equation 2. In the case of metaphorical similarity statements, the basis of comparison is usually called the *ground* (of the metaphor). This situation implies that from the point of view of someone attempting to understand a similarity statement, success can only be achieved if the intersection term is not empty. From this perspective, comprehension will be straightforward if the matching attributes are already part of the schemata for both the a term and the b term. But this need not always be the case. It sometimes happens that people have to make inferences about what a particular statement

means; that is, they sometimes have to introduce new attributes into their representation for the a term. This section focuses on the differences between such attribute-introducing statements and the more straightforward cases in which it will be supposed that the preexisting matching attributes have their salience levels in the a term promoted.

Suppose that someone utters Statement 11 with the intention of imputing sexual inertness to poor Sally. The distinction of concern is reflected in the kinds of reply that the listener might make: "Yes, I know she is. It's a shame, isn't it?" as against "Oh really? I didn't know that. I would never have thought it."

So, when we say *"a is like b,"* we may be inviting the listener not to find a match of attributes but to take some salient attributes of b that were not previously part of his or her schema for a and to build them into it.

It should be noted at the outset that strictly speaking, the difference between attribute-promoting and attribute-introducing similarity statements is not a difference that pertains to the similarity statement per se, but rather a difference that has to do with the *use* of a statement by a particular person, to a particular person, on a particular occasion. Nevertheless, given that the determination of matches plays such a central role in the present proposals, the distinction gives rise to an important question: What sort of comprehension mechanism can account for the fact that people are able to understand similarity statements when there are no matching attributes?

One answer is that the mechanism is that of attempted predication. This means that the comprehension process might be something like this: With preference being given to the most salient subschemata associated with the b term, an attempt is made to apply them to the schema for the a term. Now, one way to determine whether a particular attribute can be applied to something is to determine whether or not it is already included in the representation of that thing. Thus, matching, or at least testing for a match, might well be the first step in the process. However, matching cannot possibly always be the only step, because if the test

for a match fails, it is not possible to conclude that the attribute in question cannot be applied but only that it is not already present. The simplest prospect if the match test fails would be to determine whether any gross conceptual incompatibility would result by applying the attribute in question to the concept. This incompatibility would have to be unresolvable. Thus, for example, the attribute "being white" cannot be applied to the concept *sermons* as part of the process of trying to apply high-salient attributes of *sleeping pills* to the concept *sermons*, because to do so would result in an unresolvable incompatibility. White things are physical objects; sermons, although often sufficiently boring to induce sleep, are not physical objects.

The results of comprehension vis-à-vis changes in the nature of the schema for the *a* term enable us to characterize the difference between the two types of similarity statements. In the case of attribute-promoting statements, subschemata corresponding to the matching attributes already exist in the schema for *a*, and all that happens is that their salience levels are increased. In the case of attribute-introducing statements, either the existing value of a variable is changed or a value is provided where previously there was only a variable. Furthermore, in both kinds of statements it could be assumed that the attention to a particular variable (or to a set of them) that is occasioned by the use of the statement gives rise (at least temporarily) to an increase in the salience level of that variable (or set of them) for the *a* term. This kind of account permits the comprehension of both kinds of similarity statements, be they literal or nonliteral, to be absorbed into the general framework of comprehension that is suggested by schema theory. It has the added advantage of suggesting that from the point of view of comprehension, there may not be much difference between similes and their corresponding metaphors: Both proceed by attempted prediction. In addition, it is consistent with the view that nonliteral similarity statements do not require a different comprehension mechanism from literal ones (Ortony, Schallert, Reynolds, & Antos, 1978).

What happens if no high-salient attributes from B seem to apply to the *a* term? How can one account for the apparent fact that one can usually dream up a context in which even the most anomalous statements can make sense? One answer to this question, already hinted at, is that a speaker (or listener, reader, or writer) may reorder the salience of the attributes of (especially) the second term in the comparison. Such reordering is clearly involved when the kind of foregrounding described with respect to Statement 13 occurs. Foregrounding, it may be recalled, promotes the salience of some attribute or group of attributes, often with the result of reducing the degree of metaphoricity. Thus Statement 3 can be made interpretable by presupposing a context in which being a physical object is very important. So, a context in which it makes sense to utter Statement 18 will also be able to support an interpretation of Statement 3:

3. Billboards are like pears,

18. insofar as they are both physical objects, billboards are like pears.

However, even though Statements 3 and 18 are perfectly interpretable in a context that permits the reordering of attributes, without such an assumption they remain very odd.

The basic proposal, then, is that similarity statements are processed by attempted predication. It may be, however, that if the process fails to find a match of high-salient attributes, or more specifically, if it fails to find a literal interpretation, then the matches that are subsequently found come to be perceived as being more important than they otherwise would. For example, the simile (Statement 19, below) seems to have a higher degree of subjective similarity (if understood) than would be predicted by Equation 1, even allowing for the proposed modifications (Equation 2):

19. Cigarettes are like time bombs.

Furthermore, it appears that similes are rather like jokes in the sense that if an initial failure to properly comprehend is later followed by a full understanding, the anomalous components lose their force.

In a pilot study designed to investigate this question, subjects rated the similarity of the terms in similes as consistently higher when encouraged to perceive their metaphorical nature than under conditions that encouraged literal interpretations. Lower ratings in the literal condition were found for 14 out of 15 items. This suggests that people may reduce, perhaps to zero, the weights accorded to distinctive attributes (α and β) on discovering that they are working with a nonliteral comparison.

The central point of this section has been to offer an explanation of how people can understand comparisons when they have insufficient information about the subject term to enable them to find matching attributes. It was suggested that a mechanism of attempted predication of attributes of the b term to the a term could handle this problem. Testing for a match would then be one way to determine the applicability of an attribute. Some effects of context on salience levels were also discussed, particularly those that by permitting reordering of salience levels enable a comprehensible interpretation to be imposed on a superficially anomalous similarity statement.

Metaphor

The present proposals have a number of implications for a theory of metaphor—implications that seem to raise serious problems for some widespread assumptions while throwing light on others.

The idea that metaphors are really just covert comparisons, or analogies, has dominated philosophical approaches to the problem for centuries and has begun to take hold in cognitive psychology and psycholinguistics. This view, which is the standard Aristotelian one, has recently been proposed as the basis for a psychological treatment of metaphor comprehension by, for example, Kintsch (1974) and Miller (in press). In many ways it is an appealing notion because it not only offers the prospect of accounting for a difficult concept like metaphor in terms of an easier concept like comparison, but it also recognizes that comparisons do underlie metaphors, even though this may not be apparent from surface structural features. The present proposals, however, are incompatible with this view if it is taken to be a philosophical or psychological explanation.

According to the account that I have developed, metaphoricity is a characteristic of similarity statements. So, the similarities that are covert in metaphors are themselves metaphorical in nature. It follows from this that nothing can be explained about metaphoricity itself by (correctly) identifying a metaphor's relationship to its corresponding simile (if it has one). If the metaphor,

20. encyclopedias are gold mines,

is a meaningful statement by virtue of the fact that encyclopedias are *like* gold mines, that does not explain what it is about Statement 20 that makes it a metaphor, since encyclopedias are not really like gold mines at all, they are only metaphorically like them. So, although it is perfectly true that there may be little difference between a metaphor and its corresponding simile (indeed, if the comprehension mechanism is one of selective attribute prediction,[2] it could be the same for both), any problems about the nature of metaphors and the comprehension mechanisms for them remain unsolved.

The mistake in the view here being criticized is to suppose that similes are "literal" comparisons. This supposition is implicit in the view of Kintsch (1974) that (necessarily semantically anomalous) metaphors are reinterpreted as semantically acceptable (explicit) comparisons, and in the view of Miller (in press) that the (only) difference between a simile (to which a metaphor can be reduced) and a literal comparison is that in the latter the basis of the comparison is "obvious." It is quite explicit in Searle (in press), who refers to similes as "literal similes."

The view that the comparisons underlying

---

[2] It is possible that there are cases in which, or occasions on which, attributes of the b term are predicated, not by selecting applicable, highly salient ones, but by rejecting inapplicable ones and applying the rest. Such application by attribute rejection rather than by attribute selection might result in richer representations in the modified topic schema.

metaphors are unproblematic literal comparisons can be shown to have problems independently of the current proposals. There is linguistic evidence from the use of certain kinds of hedges (see Lakoff, 1972). Consider the hedges "metaphorically speaking" and "really." If two things are really similar, it should seem odd to say that they are metaphorically similar, and, conversely, if they are metaphorically similar, it should seem odd to say that they are really similar. Statements 21 and 22 support these predictions:

21. Metaphorically speaking, education is like a stairway,

22. education is really like a stairway.

Whereas Statement 21 makes perfectly good sense (some people might even think that it is true), Statement 22 is false. If something is genuinely metaphorically like something else, it cannot be really (literally) like it. The converse is that if two things are really similar, it makes no sense to say that they are metaphorically similar. So, although Statement 23 seems reasonable, Statement 24 does not:

23. Stairways are really like escalators,

24. metaphorically speaking, stairways are like escalators.

Similarly, one would expect that the negation of a nonliteral similarity statement interpreted literally would be true, whereas the negation of a literal one would be false, as in Statements 25 and 26:

25. Science is not (really) like a glacier,

26. an icefall is not (really) like a glacier.

Linguistic data like these support the idea that the locus of metaphoricity lies not in the surface structure of a statement (e.g., the presence or absence of "like") but in the underlying comparison itself. The standard view of metaphor as covert comparison has no way of accounting for these observations, whereas the proposals for nonliteral similarity that have been presented actually predict them.

If the inadequacy of the standard view is accepted, it follows that a second widespread assumption about the nature of metaphors suffers from the same problem. This is the view that metaphors are based on or are covert analogies. Again, the problem is not that this observation might be false. The problem is that even if true, the claim has little or no explanatory value. Analogies are statements of similarity—similarity between relations between "objects," rather than between objects themselves, but none the poorer for that. If metaphoricity is a characteristic of similarity statements, one would expect analogies to vary in their degree of metaphoricity in just the same way as simple similarity statements do. Consider, for example, Statements 27 and 28:

27. Lambs are to sheep as kittens are to cats,

28. putting William Miller in charge of the Federal Reserve Board would be like putting Rudolph Nureyev in the ring against Leon Spinks for the heavyweight title of the world.

Both are analogies, but Statement 27 is "really" true, whereas Statement 28 (William Proxmire's remark at Miller's senate confirmation hearings [3]) is not.

Thus, it can be argued that Statement 28 is a metaphorical analogy. There are no high-salient attributes (i.e., relations) of the second term that are of comparably high salience for the first (viewing the terms as the set of relationships between each pair of constituents). Statement 27 is a literal analogy because there are shared high-salient attributes of both. The pairs of constituents in an analogy can be regarded as complex schemata. The attributes of such complex schemata are the sets of relations of which they are constituted. An important difference between statements of (analogical) similarity and simple similarity statements is that in analogies, the complex schemata

---

[3] It is interesting to note the radical ambiguity of Proxmire's assertion. It could either be interpreted in terms of the high salience of Nureyev's potential nimbleness and elusiveness or, antithetically, in terms of the high salience of Spinks's superior strength and power. It takes only a little imagination to decide which of these alternatives Proxmire had in mind.

that are produced by the two pairs of terms usually have to be constructed at the time of comprehension rather than retrieved from memory as already available schemata. This means that one would expect metaphors based on underlying analogies (often called "proportional" metaphors) to be more difficult to understand. There is some evidence that this is true developmentally (Billow, 1975).

Metaphorical analogies are not at all uncommon in ordinary discourse. For example, the following analogy is taken from a story ("Your Cover is Showing," 1977) about the alias program for reinstating safe lives for informers. The story opens with the following analogy:

29. Informers are to criminal justice what uranium is to a nuclear reactor . . . .

Left unexplained, Statement 29 is somewhat obscure, perhaps because in constructing a complex concept for the second term ("what uranium is to a nuclear reactor"), all kinds of relationships can be introduced; uranium is the fuel, it is one of the more dangerous aspects of a reactor, and so forth. Yet none of these high-salient relations are obvious, high-salient relations of the first term (what "informers are to criminal justice"). A literal analogy would have a match of high-to-high salient relations; this one does not. Thus, one is willing to deny that it is true, literally, just as one is willing to deny that non-literal statements of similarity in general are true. And, just as with similarity statements in general, the explicit statement of the basis of the comparison that follows serves to reduce the perceived metaphoricity by enhancing the salience of a particular attribute. The second part is essential if the entire sentence is to be understood: "Informers are to criminal justice what uranium is to a nuclear reactor—they make the system go, but they're an awful lot of trouble to dispose of afterward."

Interesting things happen to analogies when their terms are omitted. For example, one can convert Statement 29 into Statement 30 in a way that makes it even more obscure. Now the missing term has to be supplied, but knowing what it is seems to presuppose knowledge of the yet-to-be-established relation.

30. Informers are the uranium of criminal justice.

It is interesting to note, in this example, that the most natural interpretation is far removed from that for Statement 29. Now it seems that uranium is functioning to highlight attributes related to value and scarcity, so that the most natural interpretation is that informers are very valuable to criminal justice. This confirms the claim made earlier, that attributes, since they can be complex, can often be equivalent to relations, even though they may look like simple predicates. However, to express relations in nonrelational ways can, as in this case, be very misleading. Since no sharp distinction is being made between statements of similarity that are fundamentally analogical in character and those that are not, this matters little. But it is not very encouraging for those (e.g., Miller, in press) who would argue for a conversion process of metaphors to similes and analogies as the essence of the underlying comprehension mechanism. Thus, the Aristotelian notion that metaphors are based on the principles of analogy is not very helpful. Neglecting for the moment the distinction between metaphors and similes, it has to be concluded that metaphors, like analogies, are based on the principles of similarity.

Even if it is true that there is no fundamental difference between a simple statement of similarity and an analogy, it does not mean that a theory of similarity judgments is ipso facto a theory of the problem solving that goes into the solution of analogy problems. The present proposals have nothing to say about the manner in which the complex concepts are constructed. In standard analogy problems, part of the problem is to construct a schema that involves one of the pairs of concepts in some central way in such a manner that the relation between them can be applied to the other side of the "equation." The approach to similarity being advocated here is neither capable of nor intended to deal with the way in which analogy problems are solved.

It is now possible to explicate the relationships between (a) metaphors and (b) similes and analogies. Essentially, a metaphor is a nonliteral comparison either between objects or between relations between objects. In the former case it is related to a nonliteral similarity statement that is normally called a simile; in the latter case it is related to a nonliteral similarity statement that is normally called an analogy. In both cases, the difference between the metaphor and its corresponding similarity statement is not that one is metaphorical and the other literal; the difference is that one is an *indirect* statement whereas the other is a *direct* one. Loosely stated, metaphors are indirect speech acts; similes are not. The fundamental mistake in the standard view is that it confuses metaphoricity with indirection.

Finding the comparison view of metaphor unsatisfactory, there have been those (most notably, Black, 1962, in press) who have proposed an alternative account of metaphor known as the interaction theory. The idea behind this view is that the two terms in a metaphor somehow interact to produce some new emergent meaning. The view has been criticized on the grounds that the notion of interaction is too vague and itself too metaphorical (e.g., Black, in press; Searle, in press). Although equally vague about how to characterize interaction, Verbrugge and McCarrell (1977) found evidence leading them to conclude that it is not sufficient to assume that "the topic is 'passively' schematized by salient properties of a vehicle domain: The topic and the vehicle terms *interact* in specifying the ground (p. 528)." According to the present proposals, the ground of a metaphor will be the shared attributes of the underlying nonliteral similarity statement and, in particular, those attributes that are of high salience for the vehicle but of low salience for the topic. Consequently, it would seem that a notion of interaction could be captured by the fact that the attributes comprising the ground will depend on both the topic and the vehicle. The topic and vehicle interact in the sense that the topic term imposes constraints on the attributes of the vehicle term that can be applied and that

are of high salience for the former but of low salience for the latter.

If the account of nonliteral similarity that has been presented is correct, it has some particularly important consequences developmentally. It is common knowledge that children appear to produce metaphors long before there is evidence that they can understand them, a state of affairs that if true would be completely contrary to what is generally believed about language acquisition, namely, that comprehension precedes production. The point was well put by Winner, Rosenstiel, and Gardner (1976), who started their article with the following paragraph:

> There is an apparent paradox concerning the development of metaphoric sensitivity in children. It has been argued, on the one hand, that the capacity to understand metaphoric figures of speech develops only during late childhood and early adolescence (Asch & Nerlove, 1960; Elkind, 1969; Schaffer, 1930). This contention is consistent with the view that metalinguistic skills (the ability to perform operations on language itself) develop only in the final stages of language acquisition (Inhelder & Piaget, 1958; Jakobson, 1960 . . .). On the other hand, studies focusing on the child's ability to produce figurative language have repeatedly documented the spontaneous use of metaphors, similes, and other figures of speech by preschool-age children (Carlson & Anisfeld, 1969; Chukovsky, 1968; Gardner, 1973; Gardner, Kircher, Winner, & Perkins, 1975; Weir, 1962). In addition, preschoolers have been shown to be able to match words to elements from other sensory modalities in a metaphor-type paradigm (Gardner, 1974). (Winner et al., 1976, p. 289)

Winner et al. (1976) attempted to resolve the paradox by distinguishing between various levels of metaphoric comprehension, attempting to show that very young children do have some rudimentary forms of metaphoric comprehension after all.

If the present proposals are accepted, however, there is an alternative way of resolving the paradox—a way that is perhaps more appealing. It could be argued that preschool children who are apparently producing metaphors are instead producing statements based on underlying literal similarities *from their perspectives*. For example, Chamberlain and Chamberlain (1904) cited the case of a child who used the word *moon* to refer to cakes (among other things). Now, it would of

course be possible to argue that since cakes are not really moons, the child was speaking metaphorically, but that hardly seems likely. A more reasonable approach would be to deny that the child was speaking metaphorically, asserting that very young children lack the metalinguistic awareness needed. One would then reason that from the child's perspective, the moon and cakes shared a high-salient attribute, that of being roughly circular in shape. Thus, moons and cakes were literally similar given the nascent state of the child's schemata. Later, as those schemata developed into ones more closely approximating an adult's, the matching attribute would be high in salience for the b term and low in salience for the a term, thus satisfying the criterion for nonliteral similarity. Consequently, from the adult's perspective, productions of children may appear to be metaphorical because the adult has substantially differently structured schemata and (presumably) more highly differentiated salience information for their components. The child, on the other hand, perceives the two things in question as being really very similar. The resolution of the paradox would thus depend on the recognition that the knowledge representations of adults and children are often substantially different.

Finally, one might ask how the present proposals relate to the goodness of metaphors. There has been all too little research into this question, although some is beginning to emerge. Recent research by Sternberg and his colleagues (Sternberg, Tourangeau, & Nigro, in press; Tourangeau & Sternberg, Note 1, Note 2) is based on a geometric approach. Their research assumes a representation in which the topic and the vehicle of the metaphor (the first and second terms, respectively) are viewed as belonging to different subspaces within a more global hyperspace. When the two terms are juxtaposed in a metaphor, one can think of superimposing their corresponding local spaces, coordinating the dimensions. Then, the goodness of a metaphor can be characterized in terms of the within-subspace distance and the between-subspaces distance. If the local subspaces from which the two terms are drawn are remote but the superimposed within-sub-

space distance is small, then we have a good metaphor. They also hypothesized that the aptness and the comprehensibility of a metaphor goes down if the between-subspaces distance gets exceedingly large. In many respects, the account offered by Sternberg and his colleagues is compatible with the views that have been presented here.

The present proposals are that metaphoricity requires high-salient attributes of the vehicle to be (applicable as) low-salient attributes of the topic while there exist high-salient attributes of the vehicle that cannot be applied to the topic at all. This latter constraint can be interpreted as requiring the domains (local subspaces) from which the terms are drawn to be different. Although their theory does not demand salience imbalance as a condition for metaphoricity, Tourangeau and Sternberg (Note 2) found empirical evidence that vehicles (but not topics) that were extreme on relevant dimensions within their domains tended to produce more comprehensible metaphors. They attribute this finding to the possibility that "the extreme values in the vehicle help us find its relevant characteristics by making them more salient" (p. 37). Since the high correlation between extremity and comprehensibility was not found with topics of metaphors, their results can be interpreted as evidence of high B/low A matches in similarity statements that were understood as metaphors.

The present proposals have not addressed the question of the quality or aptness of metaphors; however, it seems reasonable to suppose that quality will be primarily dependent on metaphoricity. If both metaphoricity and similarity are low, the similarity statement will appear to be anomalous and will be correspondingly difficult to comprehend. Apart from this, no special relationship between goodness and comprehensibility is predicted. This account, in fact, gains some support from Tourangeau and Sternberg (Note 2), who failed to find their predicted negative relationship between comprehensibility and between-subspaces distance but who did find a positive relationship between comprehensibility and aptness.

The question of the goodness of metaphors is a much more complex one than might be supposed at first glance. Kintsch (1974), for example, pointed out that the phrase *bachelor girl* seems much better than the phrase *spinster boy,* assuming that both are intended to be metaphorical expressions. If these expressions are cast into similes ("Some girl [or other] is like a batchelor" and "Some boy [or other] is like a spinster"), it would be necessary to show how the degree of metaphoricity of the former was higher than that of the latter. The only hope would be to find that the relative difference between the levels of salience of shared attributes was higher in the one case than in the other, and/or that a number of attributes (perhaps emotive ones) were present in the schema for one of the concepts (e.g., *spinster*), that were not present in the schema for the other. Thus, for example, attributes such as "straightlaced" and perhaps "prudish" might be of high salience for *spinster* (and consequently, "unmarried" might be of relatively lower salience), whereas these attributes might not exist in the schema for *bachelor* at all (and probably do not). Whether or not this is the correct account for these examples, it is clear that the explanation could be along these lines.

Conclusion

For 2,000 years, scholars have believed that similarity and analogy are heavily implicated in metaphors, yet there has been no satisfactory statement of the exact relation between these concepts. This article has proposed an account of similarity that offers the prospect of determining just how similarity, analogy, and metaphoricity relate to one another.

It has been argued that the essence of metaphoricity is salience imbalance. This imbalance can be enhanced by attribute inequality. An account of similarity incorporating this notion was proposed, and its implications for the asymmetry of similarity statements were explored. The proposals predict that nonliteral similarity statements will tend to be much less reversible than literal similarity statements and that in cases in which reversals still result in meaningful comparisons, the meaning change will be greater for similes than for literals. The account also predicts that the terms in similes will be perceived as being more similar than theories based on literal similarity alone would predict. Whereas the full consequences of the proposals for theoretical and empirical research on metaphors will have to await further work, a number of suggestions in this direction were made.

It should be reiterated that metaphoricity is a characteristic of similarity statements (and judgments). This observation alone could have important consequences for the way in which similarity is conceptualized and for the kind of empirical research that is undertaken in the area of metaphor. We cannot hope to properly understand the contribution that the perception of similarity makes to cognition if we fail to recognize that people are not restricted to judging, perceiving, and talking about literal similarity. Much of cognition depends on the recognition of metaphorical relations—on going beyond literal similarity.

Reference Notes

1. Tourangeau, R., & Sternberg, R. J. *Understanding and appreciating metaphors* (Tech. Rep. 11). New Haven, Conn.: Yale University, Department of Psychology, June 1978.

2. Tourangeau, R., & Sternberg, R. J. *Aptness in metaphor* (Tech. Rep. 13). New Haven, Conn.: Yale University, Department of Psychology, October 1978.

References

Anderson, R. C., & Ortony, A. On putting apples into bottles: A problem of polysemy. *Cognitive Psychology,* 1975, 7, 167–180.

Asch, S. E., & Nerlove, H. The development of double function terms in children. In B. Kaplan & S. Wapner (Eds.), *Perspectives in psychological theory.* New York: International Universities Press, 1960.

Barclay, J. R., Bransford, J. D., Franks, J. J., McCarrell, N. S., & Nitsch, K. E. Comprehension and semantic flexibility. *Journal of Verbal Learning and Verbal Behavior,* 1974, 13, 471–481.

Berlin, B., Breedlove, D. E., & Raven, P. H. General principles of classification and nomenclature in folk biology. *American Anthropologist,* 1973, 75, 214–242.

Billow, R. M. A cognitive developmental study of metaphor comprehension. *Developmental Psychology*, 1975, *11*, 415–423.

Black, M. Metaphor. In M. Black, *Models and metaphors*. Ithaca, N.Y.: Cornell University Press, 1962.

Black, M. More about metaphor. In A. Ortony (Ed.), *Metaphor and thought*. Cambridge, England: Cambridge University Press, in press.

Carlson, P., & Anisfeld, M. Some observations on the linguistic competence of a two-year old child. *Child Development*, 1969, *40*, 565–575.

Chamberlain, A. F., & Chamberlain, J. C. Studies of a child. *Pedagogical Seminary*, 1904, *11*, 264–291.

Chukovsky, K. *From two to five*. Berkeley: University of California Press, 1968.

Elkind, D. Piagetian and psychometric conceptions of intelligence. *Harvard Educational Review*, 1969, *39*, 319–337.

Gardner, H. *The arts and human development*. New York: Wiley, 1973.

Gardner, H. Metaphors and modalities: How children project polar adjectives onto diverse domains. *Child Development*, 1974, *45*, 84–91.

Gardner, H., Kircher, M., Winner, E., & Perkins, D. Children's metaphoric productions and preferences. *Journal of Child Language*, 1975, *2*, 125–141.

Grice, H. P. Logic and conversation. In P. Cole & J. L. Morgan (Eds.), *Syntax and semantics* (Vol. 3): *Speech acts*. New York: Academic Press, 1975.

Inhelder, B., & Piaget, J. *The growth of logical thinking from childhood to adolescence*. New York: Basic Books, 1958.

Jackendoff, R. A system of semantic primitives. In R. Schank & B. L. Nash-Webber (Eds.), *Papers from the conference on theoretical issues in natural language processing* (TINLAP-1). Cambridge, Mass.: Bolt, Beraneck, & Newman, June 1975.

Jackendoff, R. Toward an explanatory semantic representation. *Linguistic Inquiry*, 1976, *7*, 89–150.

Jakobson, R. Linguistics and poetics. In T. Sebeok (Ed.), *Style in language*. Cambridge, Mass.: M.I.T. Press, 1960.

Kintsch, W. *The representation of meaning in memory*. Hillsdale, N.J.: Erlbaum, 1974.

Krumhansl, C. L. Concerning the applicability of geometric models to similarity data: The interrelationship between similarity and spatial density. *Psychological Review*, 1978, *85*, 445–463.

Lakoff, G. Hedges: A study in meaning criteria and the logic of fuzzy concepts. In *Papers from the Eighth Regional Meeting, Chicago Linguistics Society*. Chicago: University of Chicago Linguistics Department, 1972.

Miller, G. A. Images and models, similes and metaphors. In A. Ortony (Ed.), *Metaphor and thought*. Cambridge, England: Cambridge University Press, in press.

Minsky, M. A framework for representing knowledge. In P. H. Winston (Ed.), *The psychology of computer vision*. New York: McGraw-Hill, 1975.

Ortony, A., Reynolds, R. E., & Arter, J. A. Metaphor: Theoretical and empirical research. *Psychological Bulletin*, 1978, *85*, 919–943.

Ortony, A., Schallert, D. L., Reynolds, R. E., & Antos, S. J. Interpreting metaphors and idioms: Some effects of context on comprehension. *Journal of Verbal Learning and Verbal Behavior*, 1978, *17*, 465–477.

Rips, L. J., Shoben, E. J., & Smith, E. E. Semantic distance and the verification of semantic relations. *Journal of Verbal Learning and Verbal Behavior*, 1973, *12*, 1–20.

Rosch, E. Cognitive representations of semantic categories. *Journal of Experimental Psychology: General*, 1975, *104*, 192–233.

Rosch, E., Mervis, C. B., Grey, W. D., Johnson, D. M., & Boyes-Braem, P. Basic objects in natural categories. *Cognitive Psychology*, 1976, *8*, 382–439.

Rumelhart, D. E., & Ortony, A. The representation of knowledge in memory. In R. C. Anderson, R. J. Spiro, & W. E. Montague (Eds.), *Schooling and the acquisition of knowledge*. Hillsdale, N.J.: Erlbaum, 1977.

Schaffer, L. F. Children's interpretation of cartoons. *Contributions to Education*. Columbia Teachers College, 1930, *429* (Whole).

Schank, R. C., & Abelson, R. P. *Scripts, plans, goals and understanding*. Hillsdale, N.J.: Erlbaum, 1977.

Searle, J. Metaphor. In A. Ortony (Ed.), *Metaphor and thought*. Cambridge, England: Cambridge University Press, in press.

Sternberg, R. J., Tourangeau, R., & Nigro, G. Metaphor, induction, and social policy: The convergence of macroscopic and microscopic views. In A. Ortony (Ed.), *Metaphor and thought*. Cambridge, England: Cambridge University Press, in press.

Tversky, A. Features of similarity. *Psychological Review*, 1977, *84*, 327–352.

Verbrugge, R. R., & McCarrell, N. S. Metaphoric comprehension: Studies in reminding and resembling. *Cognitive Psychology*, 1977, *9*, 494–533.

Weir, R. *Language in the crib*. The Hague, The Netherlands: Mouton, 1962.

Winner, E., Rosenstiel, A. K., & Gardner, H. The development of metaphoric understanding. *Developmental Psychology*, 1976, *12*, 289–297.

Your cover is showing. *Newsweek*, November 28, 1977, pp. 66–67.