EMPIRICAL STUDIES OF THE ARTS

Executive Editor
Colin Martindale, Ph.D.

Volume 11, Number 2 — 1993

Connectivity and Metaphor Comprehension

Yeshayahu Shen

Baywood Publishing Company, Inc.
26 Austin Avenue, Box 337
Amityville, NY 11701
Call (516) 691-1270 Fax (516) 691-1770
Orders only —**call toll-free** (800) 638-7819

CONNECTIVITY AND METAPHOR COMPREHENSION

YESHAYAHU SHEN
Tel-Aviv University

SECTION 1: INTRODUCTION: MAPPABILITY IN METAPHOR COMPREHENSION

Section 1.1: The Hierarchy of Mappability

According to a view commonly held in linguistics, psychology, and literary theory [1-2], a metaphor (the present article uses this term to refer both to metaphors and similes) can be described as a mapping of properties between two (conceptual) domains, that is, the mapping of the properties of a certain Source Domain (S-D) onto the Target Domain (T-D)

(1)
Tom is like a Tortoise.

A. "Tom is slow" [like a Tortoise]
B. "Tom has short legs" [like a Tortoise]
C. "Tom's (face?) color is brown or green" [like a Tortoise]
D. "Tom is a living creature" [like a Tortoise]

According to this view, (1) encompasses two conceptual domains: the "Human domain" (of which Tom is a member) which is the T-D and the "Tortoise domain" which is the S-D. Assigning an interpretation to the metaphor involves the mapping of properties from the S-D (such as "slowness") onto their counterparts in the T-D [compare 1-2].

Note, however, that in such a mapping not all the properties of the Source domain get mapped onto the Target domain. For example, in mapping the "Tortoise domain" onto the "Human" domain properties such as "being an animal" do not map. It is thus unlikely to find a reading such as "Tom is an

© 1993 Baywood Publishing Co., Inc.

animal" for the above metaphor. Obviously, what rules out this as well as other mappings is the constraints imposed by the Target domain: only those properties of the S-D which find a "matching" property in the T-D can be mapped. Thus a property such as "slowness" matches its counterpart in the Human domain (for a more detailed discussion of the problems involved in assuming this notion of "matching", (cf. Searle 1979 [3], Ortony 1979 [4]); whereas other properties of tortoises, such as "being an animal," do not meet this "matching" constraint and therefore do not get mapped (at least not literally) onto the Target domain. Thus one type of constraint on what gets mapped in metaphor comprehension is imposed by the specific properties of the Target domain.

Note, however, that of the properties that meet this constraint, i.e., the properties that are in principle mappable on the target domain, not all are equally or readily mappable. A case in point are the properties mentions in (1A-D). Although properties such as "slowness," "having short legs," "being brown or green," etc., may all match their counterpart properties in the Target domain, they are not equally mapped onto it, as reflected in (1).

(1) presents four possible readings for the metaphor "Tom is a Tortoise." All four readings are, in principle, accepted ones, and by implication, all four properties of the "Tortoise domain" introduced in these readings can be mapped onto the Human domain; nevertheless, they are not all mapped equally. Arguably it is rather (1A) which is preferred over (1B) as a reading for the metaphor, while the latter is preferred over (1C), and so on.[1]

As a further illustration, consider:

(2)
Cigarettes are like pacifiers.

A. "Providing oral satisfaction and soothing" or "Potentially addictive"
B. "Having the same cylindrical shape"
C. "Sharing the same beige color"

As in the former case, (2) presents three possible readings for the metaphor "Cigarettes are like pacifiers"; all four properties of the "Pacifiers domain" introduced in those readings can be mapped onto the "Cigarettes" domain. Note, however, that it is rather (2A) (which emphasizes the properties "Providing oral satisfaction and soothing" and "Potentially addictive") which is preferred over (2B) ("having a cylindrical shape"), which in turn outranks the property "having a beige color," presented in (2C).

The main question that these two examples raise is the mappability question: what are the constraints that determine this preferences in mappability? In fact this issue constitutes a major part of the issue of metaphor interpretation. Clearly these "mappability preferences" play a central role in metaphor comprehension and should thus be addressed by any comprehensive theory of metaphor comprehension. In spite of this, classical theories of metaphor [cf. e.g., 5-6] have not directly addressed this question.

More recently, however, there have been a few studies of metaphor which have addressed this issue. Let me briefly review two of these: Lakoff and Turner [7] and Gentner [2, 8]. Lakoff and Turner's study proposes an account of mappability which is based on the idea of "The Great chain of being" model. The Great chain of being is a cultural model that places kinds of beings and their properties on an hierarchical scale, with "higher" beings (e.g., human beings) and properties (e.g., thought, character) above "lower" beings (e.g., animals)) and properties (instinctual properties and behavior).

In this hierarchy each form of being has all of the attribute types lower on the hierarchy and in addition a higher-level properties which distinguish it from the former. Thus, for example, animals do not have mental and character attributes (like human beings do), but in addition to instinctual attributes they share properties with lower forms of being; thus, animals share with plants, complex objects, and natural physical objects (in that order) biological, structural, and natural physical attributes, respectively.

On the basis of the above considerations, Lakoff and Turner's proposal regarding mappability can be introduced. Their idea is that understanding metaphors is based on picking out the relevant highest-level properties of the Source domain and mapping them onto the Target domain. On this account the higher-level properties of a given form of being are more likely to get mapped than its lower level properties (namely, those properties which are shared with lower-level beings).

We can illustrate the above proposal, using our own examples presented in (1) and (2) above. Thus, the fact that "slowness" is more likely to get mapped than the other properties of the "Tortoise domain" can be attributed to the fact that it represents an instinctual property of tortoises, whereas the other properties represent lower-level properties of tortoises (such as biological or structural attributes).

Another example would be the "Cigarettes are like pacifiers" metaphor. In this case, the fact that "Providing oral satisfaction and soothing" is more likely to get mapped than the other properties comprising the "pacifier domain" can, in fact, be attributed to the former's representing a functional behavior of pacifiers which is a higher-level property than, say, the structural property of the pacifier's shape.

[1] Although there are many conventional metaphors such as "slow as a tortoise" in which the property gets mapped is already conventionalized in our language, it should be noted that this fact does not affect the argument we are developing here, for the following reason. Even for such conventional metaphors I will show that other properties of the Source domain can also be mapped onto the Target domain, and that these properties constitute a hierarchy of mapping preferences which is (at least partly) indifferent to the question of conventionality. For example, in the metaphor used in (1) only the first reading (which emphasizes slowness) is based on a conventional metaphor, whereas the other readings are not. This article will focus on the mapping preference hierarchy of all possible readings.

In addition it should be emphasized that a large portion of the metaphors to be mentioned in this article are clearly novel metaphors (such as (2)).

Note, however, that in order to fully account for the data presented in (1A-D) and (2A-C) more refined distinctions are required. Thus, for example, Lakoff and Turner's proposal cannot account for the fact that "having short legs" is more likely to get mapped than "having a brown or green color, since the former does not belong to a level higher than the latter. Or, to take the pacifier's example, it is not clear how this proposal is to account for the fact that "Potentially addictive" is more likely to get mapped than "Having a cylindrical shape" which in turn is more likely to get mapped than "having the beige color": Clearly, these properties do not represent a descending order of properties according to the Great Chain of Being.

Thus, while the Great Chain of Being might play a role in determining mappability to a certain extent, it still needs to be complemented by more refined distinctions as to the hierarchy of mappability.

The second proposal is that of Gentner's "Structure mapping theory." This theory focuses on the mechanism of metaphorical and analogical mappings, and in particular on the question of what gets mapped. According to this theory a metaphorical comparison (e.g., "cigarettes are time bombs") is a mapping from a base (or a source) domain to a target domain. Of particular importance is the notion of systematicity: people prefer to map systems of predicates linked by higher-order relational predicates (such as "is bigger than") rather than to map isolated predicates (such as "is red"). In other words, the main proposal of the structure mapping theory regarding the question of what gets mapped in metaphor comprehension is that relations, between objects or between relations, are more likely to be mapped than attributes of objects.

As will become clearer in section 2, Gentner's theory and the theory to be proposed later in this article share their emphasis on the fact that the elements comprising a given domain are not independent of each other but are interrelated, and that these relations are of crucial importance to what gets mapped in metaphorical mapping. Furthermore, some of the data mentioned in the present article partially support the structure mapping proposal in that, by and large, preference is given to relational predicates over object predicates. Thus for example, in the case of "cigarettes are like pacifiers," the two properties which get mapped first are relational properties (namely, "providing oral satisfaction and soothing" and "potentially addictive"), while the properties which are less preferred are isolated predicates (such as "cylindrical in shape").

Note, however, that Gentner's distinction between object attributes and relational attributes is silent as to what is preferentially mapped within each of these two groups of predicates; that is, Gentner's theory cannot account for the fact that among object attributes as well as among relational attributes that belong to the same level (that is, either to the relations between objects or the relations between relations), there are differences as to the likelihood that a given attribute will get mapped. Thus in our "John is a tortoise" example, Gentner's theory cannot account for the fact that among the two object attributes—"slowness" and "having a brown color"— the former is preferred over the latter. Or, to use the "cigarettes are like pacifiers" example, Gentner's theory cannot account for the fact that among the two relational attributes "providing oral satisfaction and soothing" and "being sucked on," the former is preferred over the latter. In sum, then, both Gentner and Lakoff and Turner's proposals may account for some of the data regarding mappability, but both theories still leave some data unaccounted for, which calls for a more detailed analysis as to the hierarchy of mappability.

Unlike these two proposals, the present study will take a different approach toward the question of mappability; this approach is based on proposals which have been made with respect to metaphor comprehension, especially those stemming from the direction taken by Ortony et al. [4, 9]. Ortony's main proposal is based on the notion of salience. Metaphor is claimed to consist of "salience imbalance" between the two concepts (or domains) comprising it, relative to the shared or matching property: the shared property is a relatively highly salient property of the source domain (the B term in Ortony's terms) and a property of low salience in the target domain. Thus, for example, in the metaphor "Lectures are like lullabies" the shared property ("including drowsiness and sleep") is, arguably, a high-salient property of lullabies but a low-salient property of "lectures." On the basis of this observation one might reasonably draw the conclusion that Ortony, had he directly addressed the issue of mappability, would have argued that mappability correlates with salience: the more salient a property in the source domain representation, the higher its likelihood of being mapped onto the target domain (provided that it can be mapped onto the target domain in the first place).

To illustrate this line of reasoning consider again our "tortoise" example. Intuitively speaking, the preference of "slowness" over "having short legs" can be ascribed to the former's higher salience (with respect to the "Tortoise domain") over the latter, which in turn is more salient than "having the tortoise's color," and so on. Such reasoning is fully compatible with other proposals that have been made with respect to metaphor comprehension by philosophers [6], psychologists [10-11], and so on.

Considerations of this kind underlie the present study. Thus, in the next section, we will critically examine the role of several factors traditionally involved in studies of salience in determining mappability. However, let me emphasize at the outset the following crucial methodological consideration. The main question we deal with here is what determines mappability, and not the question of what determines salience, which is a separate issue. We are interested in the notion of salience and salience determinants only to the extent to which they might help us in understanding and accounting for data regarding mappability. Thus, taking this approach should not be interpreted as an (even implicit) attempt to address the issue of salience, let alone take a stance on that issue.

(Only in the concluding section will I briefly introduce some considerations regarding the contribution of the present proposal to the more general questions of salience.)

The outline of the present study is as follows. Following Tversky [10], Ortony [4], Jackendoff [12], Johnson and Malgady [13], inter alia, I will begin with a brief critical discussion of the role played by four main factors involved in metaphorical mapping (that have traditionally discussed in studies of salience): diagnostically, intensity, prototypicality, and familiarity. I will argue that all these determinants indeed play a role in determining a certain portion of the "hierarchy" of mappability, but that such accounts are partial and cannot explain certain mapping preferences. Then, in section 2, I will propose an account in terms of connectivity, i.e., an account which takes the connections between properties comprising a given concept to be a major determinant of mappability—one which, taken together with the other determinants, can account for most of the data regarding the mappability of properties.

Section 1.2: Determinants of Mappability

Diagnosticity

Perhaps the most important determinant of mappability has been diagnosticity [10, 14-17]. On this account of the mappability of a given property is determined on the basis of its diagnostic value, that is, its ability to distinguish the category in question from other categories. Thus, for example, the relatively high mappability of the "slowness" of the Tortoise (as discussed above), is attributed to its ability to distinguish Tortoises from neighboring categories (this determinant has been proposed by both linguists, such as Lyons, and psychologists, such as Clark and Tversky).

This determinant is, clearly a fundamental one, and may account for a large portion of the data regarding mappability preferences. For example, the fact that "slowness" is preferred over "being a living creature" (hence the preference for (1A) over (1D)) is directly attributed to the higher diagnostic value of the former over the latter. But the diagnosticity account, in itself, is incapable of explaining more refined distinctions of mappability preferences—such as the fact that "slowness" is more likely to get mapped than, say, "having short legs," which in turn is preferred over "having the color brown." Obviously, there is no a priori reason to believe that the property "slowness" should have a higher diagnosticity value than does "having short legs" (that is, that the extension of "the entities which are not slow" is higher than the extension of "the entities which have non-short legs"), or that the color should be higher in diagnosticity than "having a brown color."

Another example is the "Cigarettes are like pacifiers" metaphor. As previously argued, the properties "providing oral satisfaction" and "potentially addictive" are more likely to get mapped than, say, "having a cylindrical shape." However, there is no reason to believe that in a null context the diagnostic value of the first two properties should be higher than the latter. The diagnosticity account may help explain the mappability difference between "slowness" and "being a living creature"; but for the more refined differences between (1A), (1B), and (1C), other factors must be considered.

Prototypicality

Another important factor is protoypicality [18]. Prototypicality measures the extent to which a given property is associated with the prototypical members of a given concept or category, and is based on the "family resemblance" idea. The idea here is that concepts are represented in memory by "prototypical members," which are those members which share the maximal number of properties with other members within the same category. On this account the prototypicality of a given property relative to a given category contributes to the property's mappability. Properties such as "slowness," "having a shell," and "having short legs" are more prototypical and hence, so the argument goes, more likely to get mapped than, say, "heaviness," since the prototypical tortoise is not necessarily heavy. As this example shows, mappability preferences based on this "prototypicality hierarchy" may account for a certain portion of the above data. But prototypicality fares no better than diagnosticity in explaining the refined distinctions in mappability preferences that were noted earlier, such as the fact that "slowness" is preferred over "having short legs": both properties, after all, are prototypical properties of tortoises.

The prototypicality approach may account for the fact that prototypical properties are more likely to get mapped than less prototypical ones; but a fuller account must be proposed to handle the more refined distinctions.

Intensity

Yet another factor that has traditionally been discussed in this context is intensity [see 10, 19; also Jackendoff's [12] notions of centrality conditions]. The idea here is that many categories consist of graded properties, such as the intensity of color, speed, etc. The mappability of a given property is determined, according to this account, on the basis of its intensity value: the higher this value, the higher the mappability. Thus, the "slowness" of the Tortoise and the "redness" of the Apple are considered more mappable than other properties in their respective domains due to the fact that both are located at some focal point on the redness dimension and relatively high point on the speed dimensions, respectively.

This account of mappability preferences is clearly restricted in its generality and explanatory power: it can only apply to graded properties, not to non-graded ones. Thus, even if this account were an adequate one within its domain of applicability, other accounts would be required to handle the degree of

Tortoises, or "having a cylindrical shape" in the case of pacifiers. But in fact there are counter-examples which suggest that, even for graded categories, this account cannot be applied automatically. Thus one can easily think of certain focal colors that are not the most (or even one of the most) mappable property(ies) of the category in question [20]. A case in point is the "green color" of a turtle which, intuitively, occupies a focal point on the "greenness" dimension, but is judged as relatively low in its likelihood to get mapped in the case of the Tortoise example.

Familiarity

According to proponents of the familiarity account [see, e.g., 21], some of the above examples can be handled by correlating mappability with familiarity, the latter being referred to the frequency with which a certain property is associated with a certain category. On this account, the fact that the "slowness" of the Tortoise is its most mappable property is to be explained simply by the frequency with which tortoises are associated with that property. Thus slowness is the most frequently mentioned property of tortoises in fairy tales and fables; indeed, it is encoded in the language in collocations like "slow as a tortoise." These frequent associations yield a high degree of familiarity, which is the main source of mappability.

As far as mappability preferences go, familiarity no doubt plays an important role in determining the likelihood of a certain property to get mapped [see, e.g., 21] and it may account for some of the data presented so far, such as the fact that "slowness" exceeds "having short legs" in its likelihood of getting mapped. However, as in the case of diagnosticity, this parameter cannot account for the more refined cases which call upon subtler distinctions. For example, although the properties "brown (or green) color" and "having short legs" are similarly associated with our conception of a tortoise (at least we have no reason to believe that one of them exceeds the other in familiarity), they are not equally mapped on their respective target domains.

Note, moreover, that although familiarity can be used as a measure of mappability it cannot be considered as a genuine explanation for the mappability data; attributing high mappability to a high degree of familiarity does not constitute an explanation of the phenomenon in question, since the question still remains as to why this particular property, rather than other properties, is more frequently associated with the concept in question. Considering the tortoise example again, the familiarity account would not be able to explain the fact that, among other properties of tortoises, it is "slowness" rather than, say, the tortoises' color which has always been more frequently associated with tortoises. Moreover, this account would not be able to account for the fact that among the various slow-moving creatures (which also include various types of insects such as worms, bugs, and so on) the property of "slowness" has always been more frequently associated with the tortoise than with most other slow creatures.[2]

SECTION 2: THE CONNECTIVITY ACCOUNT

Section 2.1: Connectivity

Levels of Approximation in Describing Metaphorical Mappings

Having reviewed the above factors and their roles in determining mappability preferences, we may conclude that these (partially overlapping) determinants provide only a partial explanation of the data under consideration—that is, they provide us with a first approximation toward a more comprehensive account. The present section suggests that a main reason for this relative lack of success lies in the failure of these accounts to notice or appreciate the importance of connections (i.e., causal or semi-causal relations) between properties in determining the likelihood of a given property to get mapped. Note that all former accounts have to do either with the relationship between a given property and its dominating category (whether on the basis of associative familiarity or of diagnosticity), or with its position on a certain external scale or dimension (as in the case of intensity). However, a given property also shares relations with the other properties comprising its dominating category. In what follows I will focus on one specific type of such relation, namely, causal relations [cf. 2, 8, 23, who discuss the crucial role played by causal relations among properties in human categorization]. It will be argued that causal relations, of precisely this type, referred to as connections, are an important determinant of property mappability and enable us to overcome at least some of the imperfections of each of the former accounts.

It must be emphasized at the outset that this "connectivity account" is not an alternative to the other accounts, but rather a complementary account, in that it takes as its input the output of the other determinants (diagnosticity, typicality, etc.) and provides us with more refined distinctions within the "mappability hierarchy."

To start with, we will first mention briefly our general assumptions as to the nature of semantic representations. The central assumption is that knowledge is represented in schema-like representations, that is, as an organized and integral

[2] The experiments conducted by Malt and Smith [22] raise another type of empirical objection to the familiarity account. They show that to a large extent "variations in typicality can exist independent of variations in familiarity, although familiarity may also play a role" (p. 222). Inter alia, they demonstrate that structural properties, such as family resemblance, can offer a better account of typicality ratings than does familiarity. A case in point is the identification of items such as "sofa" or "a chicken," which people are highly familiar with but which are nevertheless judged to be relatively non-salient members of their respective categories ("Furniture" for "Sofa"; "Bird" for "Chicken").

"whole" rather than a list of independent properties or features. The characteristics of this kind of knowledge representations are the following.

1. The meaning of a given category or domain is represented in memory as a typical "Schema" or "scene" [see 9];
2. The internal structure of this typical schema consists of certain elements (called properties) which are also shared by other categories;
3. These properties are linked to each other by various types of Connections, the most important being causality;
4. These connections (or correlations) between properties constitute a strong basis for the concept's internal coherence [see 23].

For example, the concept of a Tortoise activates the "typical Tortoise schema" which includes properties such as a certain size, a certain color (green or light brown), having a shell, slowness, etc. Moreover, some of these properties are connected to each other via causal or "semi-causal" connections; thus, for example, both the tortoise's possession of a shell and his having short legs are (according to our folk theory) causally related to his slowness. (Note that connections can be based on our folk theories about the world and need not necessarily be true from a scientific point of view.

Let us take the notion of connectivity one step further by considering the number of connections between properties. Consider again the "Tortoise" example. Here the "slowness" of the Tortoise is connected via causal relations to: 1) The size of the Tortoise; 2) Its having a shell; 3) Its clumsy shape; 4) Its having short legs. On the other hand, the last-named property is connected only to the property "slowness." We may therefore conclude that, in the context of "a Tortoise," "slowness" has a higher degree of connectivity than "having short legs." In general, then, properties differ from each other with respect to the number of connections they have.

On the basis of the above considerations, a distinction can be drawn, for those properties comprising the prototypical representation of a given concept, between connected versus non-connected properties. "Slowness" and "Having a brown color" are examples of the former and the latter, respectively.

Additionally, among the connected properties a subtler distinction can be drawn between high-connectivity properties and low-connectivity properties on the basis of the number of connections they have, illustrated by "slowness" and "having short legs," respectively.

Note that the boundary between high-connectivity and low-connectivity properties is not clearly defined, and can vary for different concepts. Nevertheless, as is the case in other fields of semantic representation and categorization, the difficulty is only with borderline cases; the other cases clearly fall within one group or the other. In the Tortoise case, for example, it is clear that "slowness" and "having a shell" belong to the high-connectivity group: the slowness of the Tortoise is connected to its having a shell, to its having short legs, to its weight, to its having a clumsy shape, etc. Its having a shell is connected to properties such as slowness, clumsy shape, weight, having a retractile head and legs, etc. By contrast, "having short legs" is a low-connectivity property: presumably, the only property to which it is connected is the tortoise's "slowness." This analysis, and the ones that follow, are supported by judgments of connectivity elicited from several native speakers of English.

As a further illustration consider the "Pacifier" concept schema. This schema consists of properties such as "providing oral satisfaction and soothing," "potentially addictive," "having a cylindrical shape," "usually used by babies," "having a small size," "typically having a beige color," etc. According to our connectivity scale, properties such as "providing oral satisfaction and soothing" and "Potentially addictive" are considered to be highly connected: the first connects up to properties such as "Potentially addictive," "being used by babies," "having a small size," "having a certain shape"; the second property ("potentially addictive") connects up to the properties "used by children," "providing oral satisfaction," etc.

By contrast, the color of pacifiers does not connect up to any of the other properties mentioned and therefore is considered a non-connected property.

The low-connectivity group consists of properties such as "having a cylindrical shape," which presumably connects up only to the pacifier's function of oral satisfaction.

In what follows I will present four sets of data supporting the connectivity account. The first set of evidence (section 2.2.1) aims at showing that connectivity indeed plays a role in metaphor comprehension and that understanders of metaphors are sensitive to it. The next two sets (section 2.2.2) will support the connectivity account as a measure of mappability; the fourth set of evidence (section 2.2.3) provides further, informal support for the role played by connectivity in metaphor comprehension involving different tasks than the former ones.

Section 2.2: Evidence for the Connectivity Account

2.2.1 The Inference-Activation Argument

The first set of evidence is intended to establish the fact that connectivity does indeed play a role in metaphor comprehension. Note first that if the connectivity account is correct, then metaphor understanders can be assumed to be sensitive to connections between properties while interpreting metaphors. One way to examine this prediction is to compare the effects of connected versus non-connected properties on the drawing of inferences in metaphor comprehension. If it turns out that the inference structure that can be set up from metaphors based on connected properties is richer than is the case for the non-connected ones, this will argue for metaphor understanders' sensitivity to connectivity. This is because there are "traces" leading from connected properties to other properties, yielding the possibility of inferencing based on probabilistic world knowledge (as opposed

to logical inferences). By contrast, non-connected properties contain no such "traces," and hence no such inferencing is possible.

In order to examine the validity of this hypothesis I conducted an informal experiment. Before introducing the experiment, however, let me briefly add a methodological note. The experiment to be described deals only with the distinction between connected and non-connected properties, and ignores low-connectivity properties. The reason is that, unlike the reading preference experiment to be presented in the next (sub)section, the present experiment is aimed at supporting the claim that metaphor understanders are sensitive to connectivity while interpreting metaphors, rather than the claim that connectivity is a reliable measure of mappability. For this reason, I was not interested in quantitative differences involving the number of connections, but rather in the qualitative differences between connected and non-connected properties, and their effects on inference structure in metaphor comprehension.

The experiment was conducted as follows. A group of six informants were presented with ten pairs of metaphorical statements, whose shared property (the ground) was explicitly mentioned. Each pair consisted of two metaphorical statements, one containing a connected property as ground, the other a non-connected property (e.g., "He is as heavy as an elephant" and "He has a long nose like an elephant," respectively). We will call the first and second metaphor in each set "connected metaphors" and "non-connected metaphors," respectively.

In order to examine the number of inferences that can be drawn from "connected" and "non-connected" metaphors, informants were presented with a set of possible inferences for each pair of metaphorical statements; these inferences contained five statements containing salient properties of the source domain (other than those mentioned in the metaphorical statement itself) that can, in principle to be mapped onto the Target domain. For example, if the metaphorical statement was "He is as heavy as an elephant," then the possible inferences would involve properties such as "tallness," "strength," "having thick legs" or "having a long nose."

For each metaphorical statement and corresponding set of inferences, the informants were then asked to rate the likelihood that each of the successive properties would be inferred from the target statement.

The instructions read as follows. "You may be aware of the fact that the inferences we usually draw in our everyday life are not necessarily logical or scientifically true. Suppose you are told, e.g., that someone is a thief; without knowing anything about this person, you are likely to infer that he belongs to a poor family, that he is not educated, etc., although obviously there is nothing logical or necessarily true about this. In the following experiment you will be asked to draw this kind of "likely-to-be-drawn" inferences from various comparisons that you will read."

After this introduction, informants were presented with the following task (illustrated here by a connected metaphor).

Suppose you are told that "John is as heavy as an elephant"; would it then be likely or unlikely to infer that:

1. John is tall.
2. John is strong.
3. John has thick legs.
4. John has a long nose.
5. John has long ears.

It was hypothesized that informants would draw more inferences from connected metaphors (e.g., "John is as heavy as an elephant") than from their non-connected counterparts (e.g., "John has a long nose like an elephant").

The results strongly support the above hypothesis. Fifty-eight of sixty judgements that were made (6 informants multiplied by 10 metaphors), revealed the pattern predicted by the connectivity account: the number of inferences informants tend to draw is significantly greater with connected (mean = 1.633) than with non-connected (mean = 0.1166) metaphors, across informants and across metaphors. For example, in the connected condition of the above metaphor pair (viz. "John is as heavy as an elephant"), informants tend to judge several inferences as likely (such as: John is tall; John is strong, etc.), whereas no such inferences were drawn from the non-connected member of the pair (viz. "John has a long nose like an elephant"). This pattern held for the other nine metaphor pairs used in the experiment.

Two tailed two sample t tests were performed, comparing the means of both conditions (connected and non-connected conditions, $N = 60$ for each condition). The results are statistically significant ($t = 10.98$, $df = 118$, $p < .05$).

These data directly support the connectivity account by showing that it plays a role in metaphor comprehension. Note, moreover, that the connectivity account is the only approach which can account for these data, for it is the only account which takes into consideration the causal (or semi-causal) relations between properties comprising a given category.

2.2.2 *Connectivity as a Measure of Mappability*

2.2.2.1 Preference of Readings for Metaphors — Having established that connectivity indeed plays a role in metaphor comprehension, let us turn to another set of evidence which support the connectivity account as providing a measure of mappability in metaphor comprehension. The connectivity account implies that, among the properties that can be mapped from the source onto the target domain and which are equally mappable according to the other determinants, high-connectivity properties will be preferred over low-connectivity properties, which in turn will be preferred over non-connected properties.

Let us consider the following examples of preferences of reading for metaphors. (3) below presents such a list of readings, ordered according to a preference scale.

(3)
Tom is like a Tortoise

A. "Tom is slow," or "Tom has a shell" [like a Tortoise];
B. "Tom has short legs" [like a Tortoise];
C. "Tom is brown or green" [like a Tortoise].

The readings in (3A), (3B), and (3C) correlate with the scale of connectivity. Recall that the TORTOISE schema consists of properties such as slowness, a certain size, short leggedness, a certain shape, a particular color (green or light brown). As argued in the previous section, the properties introduced in (3A) (namely, the "slowness" of the Tortoise and its "having a shell") are high-connectivity properties, and in (3B) ("having short legs"), low connectivity. As to (3C) (the color of the Tortoise), it is presumably lowest in connectivity among the typical properties of the Tortoise, as it does not seem to connect to any of the other properties. Hence, the connectivity scale, consisting (roughly) of three degrees of connectivity, can account for the preference of readings in (3).

As for the second metaphor we have been using, we again present the various possible readings arranged in order of preference:

(4)
Cigarettes are like pacifiers

A. "Providing oral satisfaction and soothing," or "Potentially addictive."
B. "Having the same cylindrical shape"
C. "Sharing the same beige color."

As argued previously, properties such as "providing oral satisfaction and soothing" and "potentially addictive" are considered to be highly connected properties; properties such as "having a cylindrical shape" are considered to be low-connectivity ones; and the pacifier's beige color is a non-connected property. Here too, then, the hierarchy of connectivity correlates with the preference of readings described above.

To examine this hypothesis in a more systematic way, I conducted the following informal experiment. Six informants participated, and ten different sets of metaphors were used. Each set contained a metaphor in which the ground (i.e., the shared property) was not explicitly stated (e.g., "Cigarettes are like pacifiers"), and a set of three possible readings which differed in the degree of connectivity of the metaphorical ground: one of the readings emphasized a highly-connected property (e.g., "Cigarettes are like pacifiers in that they provide oral satisfaction"), a second reading emphasized a low-connectivity property (e.g., "Cigarettes are like pacifiers in that they have a similar cylindrical shape"), while the third reading contained a non-connected property as the ground ("Cigarettes are like pacifiers in that they have the same color").

In order to check on the reliability of my judgments as to the connectivity between properties, I presented a random selection of concepts and the properties comprising them to another judge who was asked to determine, for each pair of properties belonging to the same given concept, whether they were connected (in the sense defined above) or not. The agreement between the other judge's determinations and mine was around 90 percent (most of the disagreements were resolved through discussion).

The order of these readings was randomized. The informants were then asked to rate their preference for these readings on a scale of 1 to 3. They were given the following instructions: "In what follows you will get 10 sets of metaphors and for each metaphor you will get three interpretations. My assumption is that these interpretations differ with respect to their likelihood or soundness as interpretations of the corresponding metaphor. I would like you to rate these interpretations on a scale of 1 to 3: assign the number 1 to the interpretation which seems to be the most likely interpretation, 2 to the one which is less likely, and 3 to the least likely one."

Recall that the hypothesis was that the hierarchy of connectivity would correlate with the preference of readings, namely that the HCR would exceed LCR which, in turn, would exceed the NCR. This hypothesis consists of two sub-hypotheses: 1) that HCR would be scored higher than both LCR and NCR; and 2) that LCR would score higher than NCR.

Two tailed two sample t tests was performed, examining these two (sub) hypotheses. As to the first (sub)hypothesis, the results are statistically significant ($t = 7.5$, $df = 118$, $p < .05$). The means of the first and second condition ($N = 60$ for each condition) were 0.966 and 0.034, respectively. Situations in which HCR>LCR or HCR>NCR were dummy coded "1," all remaining situations were coded as "0".

A further examination to determine data trends in the LCR and NCR categories (the second sub-hypothesis) was performed using another two tailed two sample t test. Here also, the results are statistically significant ($t = 6.14$, $df = 118$, $p < .05$). (Situations in which LCR>NCR were dummy coded "1", whereas those in which LCR<NCR were coded as "0"). The means were 0.83 and 0.167, respectively.

As can be seen the results support both these hypotheses. By and large, the pattern predicted, indeed, applied. These data directly support the connectivity account by showing that it plays a role in reading preference for metaphors. As the theoretical considerations regarding connectivity has already been discussed in detail, no further discussion will be added.

To gain further support for the connectivity account I analyzed a small sample of metaphor readings, as provided by Preminger in his *Encyclopedia of Poetry and Poetics* [23]. Preminger's explanations typically fit smoothly into our account, as exemplified in the following metaphor: "her head . . . with its anchoring calm," explained as follows: "Her head that, with its air of calm, makes you feel secure (and hopeful) as an anchor would in a ship" [24, p. 491].

The main high-connectivity property in this case is the anchor's function, which is connected to many of the properties comprising the "anchor schema." Thus the function of the anchor is causally related to its size, shape, weight, and so on. This high connectivity of the anchor's function may account for the fact that the property which the metaphor's author chose to highlight, and indeed which immediately comes to mind (as suggested by several informal intuitive responses to the metaphor), is "causing to feel secure," which is a reasonable paraphrase of the anchor's function. Analysis of similar data from the sample likewise supports the connectivity account.

2.2.2.2 Judgments of the Ground's Adequacy — Another source of evidence is judgments as to the adequacy of the ground of the metaphor: to what extent does a particular shared property seem "reasonable" as a basis (or "ground") for comparing two domains? If the connectivity account is correct, then judgments as to the adequacy of the ground should be expected to correlate with the degree of connectivity of that ground within the relevant Source domain category. If a property with high connectivity of the source domain is selected as the ground, the adequacy rating for this ground should be relatively high, in comparison to cases where a lower-connectivity property is selected. Moreover, it is further predicted that if there is more than one property of the Source domain which matches one of the Target domain's properties, preference will be given to the property with the highest connectivity.

In order to examine the validity of this prediction, I re-analyzed the results obtained by Ortony's experiment [9]. In this experiment subjects were presented with pairs of statements, each consisting of a similarity statement and its ground. For example, the subjects were presented with a similarity statement such as "Trust is like glue" and its ground, "Producing a strong and permanent bond"; they were then asked to rate, on a scale of 1 to 4 the degree to which this ground can be considered an adequate ground for the similarity statement in question (1 = least adequate; 4 = most adequate). In Ortony's experiment each ground which was rated above 2.5 was considered as an adequate one; no further analysis as to the more refined differences in mappability was performed by Ortony.

The purpose of my reanalysis of Ortony's data was to examine the differences between ground adequacies that were rated between 2.5 and 4. This was done as follows. I divided the set of above pairs into three more or less equal groups, on the basis of their (average) adequacy ratings. The highest group consisted of pairs whose adequacy ratings were between 3.33-3.75; the next group of pairs had ratings ranging between 3.00-3.32; the third group had ratings below 3.00. The first group contained fourteen pairs, the second thirteen, the third fifteen. (The reason for this inequality of numbers was simply that no precisely equal division was possible, as more than 4 items fell right at the "group boundaries" (3.33 and 3.00)).

The results of this reanalysis were as follows. In general, a high correlation was found between the adequacy scale and connectivity. Thus twelve of the fourteen grounds in the first group (the high-adequacy ratings) represented the property with the highest connectivity within the respective Source domain category; in the second group (the medium-adequacy group) the ratio was eight out of thirteen, while in the third group (the lowest adequacy ratings) this ratio decreased to three out of fifteen. For example, the adequacy ratings for "cutting deeply and causing pain" as a ground for "insults are like razors" was 3.58. And indeed, "cutting deeply and causing pain" is also the property of the "razor" category having the highest connectivity: by virtue of representing the razor's function, it is connected to the razor's shape, size, mode of use, etc. On the present account, the high adequacy rating for this property is due to its high connectivity in the Source domain category. At the other end of the scale were examples such as "Schools are like zoos," for which the adequacy rating of the ground "being disorganized and noisy" was 2.83. On the present account this is a relatively low-connectivity property among the properties comprising the category "Zoo"; at most it can be considered as one of the by-products of the main function of Zoos, which is "being places where animals are cared for"; and indeed, when the latter property was selected as a ground for the similarity statement "Farms are like zoos" (which Ortony defines as a literal similarity), it got high adequacy rating (3.75). These data clearly confirm the main prediction: adequacy ratings depend mainly on connectivity.

In addition, the second prediction is also confirmed—that if there is more than one property of the source domain which matches one of the target domain's properties, preference will be given to the property having the highest connectivity. A typical example is the simile, "Cigarettes are like pacifiers—providing oral satisfaction and soothing" (3.41). Note that in addition to the property "providing oral satisfaction and soothing" there are other properties of the "Pacifier domain" which can be mapped onto the "Cigarette domain," such as "potentially addictive," or "cylinder shape of the part being sucked" etc. The point, however, is that the property that was selected by the subjects, and got a high adequacy rating, is indeed the property with the highest connectivity in comparison to the others: the function of "providing oral satisfaction and soothing" is connected up to most of the other salient properties of the "pacifier domain," such as the pacifier's shape, its size, its addictiveness, the way it is used" etc.—which is not the case for the other properties.

Comparing similes to literals (i.e., literal comparisons) reveals something which contrasts with a prediction implicated by Ortony's theory, according to which the adequacy of the ground will be greater for similes than for literal comparisons. Contrary to this expectation it was found that the adequacy of the ground depends on connectivity rather than on the literal/metaphor distinction. Consider for example:

Literal: Farms are like zoos—being places where animals are cared for (3.75)
Simile: Schools are like zoos—being disorganized and noisy (2.83)

Simile: Trust is like glue—producing a strong and permanent bond (3.50)
Literal: Paste is like glue—being a sticky substance sold in bottles (3.08)

In both these cases, however, regardless of whether the simile or the literal comes out on top, it is the connectivity value (with respect to the source domain) of the property represented by the ground which accounts for the adequacy ratings. Thus "being places where animals are cared for" is the function of zoos, and as such it connects up with a large number of other properties of zoos, such as their inhabitants, their general structure, their spatial organization, the fact that people come to visit them, their noisiness, the function of the guards, etc. Only some of these properties are also connected to the property "being disorganized and noisy," which is, at most, a by-product of the fact that animals do live and are taken care of in zoos.

Similarly, in the "glue" case, it is the glue's function which is highest in both adequacy and connectivity, as it is connected both to perceptual properties, such as shape, substance, etc., and to functional properties, such as the procedures for using it, etc. Presumably, the property "being a sticky substance sold in bottles" is of lower connectivity.

2.2.3 Justification

Another kind of support for the connectivity account is provided by an analysis of readers' responses to more informal experiments conducted in connection with an additional phenomenon, namely, Justification.

The connectivity account, if correct, can be extended to account for another task that readers of metaphor may face, that of Justification. Suppose the reader is given the task of justifying the fact that a certain property which in general is very unlikely to get mapped does in fact get mapped in a particular case. The data to be presented below suggest that at least in some cases this justification process is based on connecting up the property in question with other properties in the same set.

Consider for example, the simile "the sun stood there . . . still as a bull," taken from a poem written by the Israeli poet Amir Gilboa. Obviously, "stillness" is not a property of the category "bull" which is likely to be mapped, in standard context. However, in attempting to construct an interpretation for this simile, the reader has to "justify" this property with respect to the category "bull." An informal analysis of readers' responses to the question "in what sense can 'stillness' be considered a mappable property of bulls?" clearly indicates that the process of justifying is based on an attempt to establish connections between this property and other properties of bulls. A typical response would connect the "stillness" of the bull to its large size and weight (which presumably cause it, according to this ad-hoc folk theory, to remain motionless), to its being a dumb creature, etc.

A related case of justification is the one in which one is asked to justify the fact that a given property is conventionally associated with a certain concept. In order to examine this process of justification I conducted an informal experiment in which three subjects were given four conventional properties and four related categories, and were asked to justify the claim that these properties are conventionally associated with the latter. (The properties were "smartness" for "owls," "bravery" for "lions," etc.) It turned out that this process of justification takes the form of connecting the property in question to other properties of the same category. Thus when one of my subjects was asked why the "owl" is considered to be one of the smartest animals, he tried to justify it by relating smartness to the owl's eyes, which are "shaped like small professor-type eyes," referring to cartoons in which owls are depicted, etc. A similar line of justification was taken with regard to the folk belief that "bravery" is a mappable property of lions. Here again subjects tended to connect this property to other properties of the lion, such as its "majestic" behavior, the fact that it does not shy away from a fight, etc.

This and similar examples of the same process suggest that the attempt to "raise" the mappability of a given property is based on establishing connections between this and other properties of the same set. By being connected to other properties, the property in question assumes a more essential part in the schema. This process is directly related to the idea [initially developed in 23] that concepts are like theories, in that (among other characteristics) there is an attempt to establish causal connections (regardless of their scientific validity) among the relations between properties comprising a given category.

The process of justification then, provides an informal source of additional evidence for the close relation between mappability and connectivity.

CONCLUSION

The main purpose of the present study has been to highlight an important issue within metaphor comprehension, namely, mappability, or: what gets mapped in metaphors. On the present view this aspect of metaphor comprehension plays a central role in the process of assigning an interpretation to metaphors, a role which has not been sufficiently appreciated in philosophical, psychological, linguistic, and literary studies of metaphor comprehension.

More specifically, this article has emphasized the central role of connectivity in determining mappability. Several determinants can plausibly be taken as responsible for the various data regarding mappability: diagnosticity, familiarity, prototypicality, and intensity. However, in order for a theory of mappability in metaphors to capture more refined distinctions between various mappings, the former determinants must be augmented with connnectivity. The "connectivity account," then, should not be considered as an alternative to the other accounts,

but rather as a complementary account, in that it takes as its input the output of the other determinants (diagnosticity, typicality, etc.) and enables us to make more refined distinctions within the mappability hierarchy.

However, the analysis of connectivity and its role in determining mappability may be seen in a broader perspective, as indirectly contributing to the study of an important notion in theories of categorization, namely Salience. Recall that our starting point in discussing mappability was the assumption that mappability is related, at least to a certain extent, to salience; this assumption led us to examine the role played by factors which have been traditionally studied by theories of salience, in determining mappability. The main finding of the present study is that these factors should be supplemented with connectivity, a factor that has not traditionally been associated with saliency. What I would like to suggest, in a speculative manner, is that the link goes both ways: not only should a theory of mappability borrow concepts from the theory of salience, but a theory of salience can also benefit from a theory of mappability. In other words it seems to me highly likely that connectivity (which has been found to be a major determinant of mappability) can be viewed as an important determinant of salience in properties comprising natural categories (of the kind examined throughout this article). Although this possibility has yet to be explored thoroughly, there are some initial indications as to its plausibility. Let me briefly mention two domains, other than natural categories, in which connectivity has been considered a major determinant of salience or importance. I am referring to Trabasso and van den Beck's studies of story comprehension [24], and to Yekovich and Walker's study of scripted concepts [25].

Trabasso and van den Beck's study of story comprehension found that connectivity plays a central role in narrative comprehension [24]. These researchers have been developing a Causal Network model: the story is represented as a causal network in which the nodes in the network represent story statements, and the arrows between the nodes indicate the causal (or temporal) relations among them. The main observation emerging from this theory is that connectivity, defined in terms of the number of direct causal connections between propositions in the story, predicts judgments of importance. Thus the greater the number of causal connections leading from a given statement to other story statements, the higher its probability of being recalled, retained in memory, and/or assigned a high level of importance.

Another area in which connectivity plays a central role in determining importance of salience is the domain of scripted concepts, such as "painting a room." As argued by Yekovich and Walker in a recent study of scripted categories, connections between concepts comprising a scripted memory structure are a central determinant of the importance of particular concepts [25]. For example, within the memory structure for "painting a room," the concept "wall" is more important (the authors use the term "central") than "water" (assuming that water is used to clean up after painting the room). The explanation is that the former includes a higher

number of connections to other concepts within the given memory structure than does the latter.

These two studies, then, as well as some of the data presented in the present article suggest that connectivity should be regarded as a central factor or principle of organization in cognition, which will be found to determine the organization of memory structures in various conceptual domains.

ACKNOWLEDGMENTS

I would like to thank Dedre Gentner, Ray Gibbs, Adele Goldberg, Chana Kronfeld, Mark Turner and Eve Sweetser for many helpful comments and valuable suggestions on an earlier draft.

REFERENCES

1. G. Lakoff and M. Johnson, *Metaphors We Live By*, University of Chicago Press, Chicago, 1980.
2. D. Gentner, Structure-Mapping: A Theoretical Framework for Analogy, *Cognitive Science, 7*, pp. 155-170, 1983.
3. J. Searle, Literal Meaning, in *Expression and Meaning*, J. Searle (ed.), Cambridge University Press, Cambridge, England, 1979.
4. A. Ortony, Beyond Literal Similarity, *Psychological Review, 86*, pp. 161-180, 1979.
5. M. C. Beardsley, *Aesthetics*, Macmillan, New York, 1958.
6. M. Black, *Models and Metaphors*, Cornell University Press, Ithaca, New York, 1962.
7. G. Lakoff and M. Turner, *More Than Cool Reason—A Field Guide To Poetic Metaphor*, Cambridge University Press, Cambridge, England, 1988.
8. D. Gentner and C. Toupin, Systematicity and Surface Similarity in the Development of Analogy, *Cognitive Science, 10*, pp. 277-300, 1986.
9. A. Ortony, Salience, Similes, and the Asymmetry of Similarity, *Journal of Memory and Language, 24*, pp. 569-594, 1985.
10. A. Tversky, Features of Similarity, *Psychological Review, 84*, pp. 327-352, 1977.
11. R. G. Malgady and M. Johnson, Measurement of Figurative Language: Semantic Feature Models of Comprehension and Appreciation, in *Cognition and Figurative Language*, R. P. Honeck and R. R. Hoffman (eds.), Erlbaum, Hillsdale, New Jersey, 1980.
12. R. Jackendoff, *Semantics and Cognition*, MIT Press, Cambridge, Massachusetts, 1983.
13. M. Johnson and R. G. Malgady, Toward a Perceptual Theory of Metaphoric Comprehension, in *Cognition and Figurative Language*, R. P. Honeck and R. R. Hoffman (eds.), Erlbaum, Hillsdale, New Jersey, pp. 259-282, 1980.
14. H. H. Clark, Word Association and Linguistic Theory, in *New Horizons in Linguistics*, J. Lyons (ed.), Penguin, Harmondsworth, England, 1970.
15. H. H. Clark and E. V. Clark, *Psychology and Language*, Harcourt Brace Jovanovich, New York, 1977.
16. J. Lyons, *Semantics 1 and 2*, Cambridge University Press, Cambridge, England, 1977.

17. R. Tzur, *The Cognitive Theory of Metaphor*, Mechearey, Yerushalaim, 1983 (in Hebrew).
18. E. Rosch and C. B. Mervis, Family Resemblances: Studies in the Internal Structure of Categories, *Cognitive Psychology, 7*, pp. 573-605, 1975.
19. A. N. Katz, Metaphoric Relationships: The Role of Feature Saliency, *Journal of Psycholinguistic Research, 11*, pp. 283-296, 1982.
20. E. Rosch, Cognitive Reference Points, *Cognitive Psychology, 7*, pp. 532-547, 1975.
21. C. Kronfeld, *Aspects of Poetic Metaphor*, Ph.D. dissertation. University of California, Berkeley, 1983.
22. B. C. Malt and E. E. Smith, The Role of Familiarity in Determining Typicality, *Memory and Cognition, 10*, pp. 69-75, 1982.
23. D. Medin and W. Wattenmaker, Category Cohesiveness, Theories, and Cognitive Archeology, in *Concepts and Conceptual Development*, U. Neisser (ed.), Cambridge University Press, Cambridge, England, pp. 25-62, 1982.
24. A. Preminger (ed.), *Encyclopedia of Poetry and Poetics*, Macmillan, New York, 1974.
25. T. Trabasso and D. van den Broek, Causal Thinking and the Representation of Narrative Events, *Journal of Memory and Language, 24*, pp. 612-630, 1985.
26. F. Yekovich and C. H. Walker, Retrieval of Scripted Concepts, *Journal of Memory and Language, 25*, pp. 627-644, 1986.

Direct reprint requests to:

Yeshayahu Shen
Department of Poetics and Comparative Literature
Tel Aviv University
Ramat-Aviv, Tel Aviv 69978
Israel