

In Press: *Metaphor and Symbol*

Metaphor comprehension:
What makes a metaphor difficult to understand?

Walter Kintsch & Anita R. Bowles
University of Colorado

Send Correspondence to:

Walter Kintsch
Department of Psychology
University of Colorado
Boulder, CO 80309-0345
wkintsch@psych.colorado.edu
303-492-8663

Abstract

Comprehension difficulty was rated for metaphors of the form Noun₁-is-a-Noun₂; in addition, participants completed frames of the form Noun₁-is-_____ with their literal interpretation of the metaphor. Metaphor comprehension was simulated with a computational model based on Latent Semantic Analysis. The model matched participants' interpretations for both easy and difficult metaphors. When interpreting easy metaphors, both the participants and the model generated highly consistent responses. When interpreting difficult metaphors, both the participants and the model generated disparate responses.

Key Words

metaphor

latent semantic analysis

predication

comprehension

metaphor comprehension

Metaphor comprehension:

What makes a metaphor difficult to understand?

There exists a considerable and convincing body of research in cognitive psychology and cognitive science that indicates that people understand metaphors in much the same way as they understand literal sentences (Cacciari & Glucksberg, 1994; Gibbs, 1994, 2001; Glucksberg, 1998). Some metaphors are easier to understand than others, but the same can be said for literal sentences. On the whole, the view that understanding metaphors is a more complex process than understanding literal sentences is not supported by this body of research. In particular, it does not appear that metaphor comprehension first involves an attempt at literal comprehension, and when that fails, a metaphoric reinterpretation. Certainly, that is sometimes the case for complex, often literary metaphors, but most ordinary metaphors encountered in common speech and writing are simply understood without any need to figure them out. Some literal sentences, too, challenge comprehension and require a certain amount of problem solving for their comprehension. But most of the time the sentences that we hear and read are understood without deliberate reasoning, whether they are metaphorical or literal.

Of course, claiming that metaphorical sentences are understood in the same way as literal sentences does not tell us how either one is understood. Here, we describe a model of text comprehension (Kintsch, 1998, 2001) that attempts to specify the process of comprehension for both literal and metaphorical sentences, simulate the computations involved, and evaluate the model empirically.

A basic assumption of this model is that the meaning of a word, sentence, or text is given by the set of relationships between it and everything else that is known. This idea is operationalized in terms of a high-dimensional semantic space. Words, sentences, and texts are represented as vectors in this space; that is, meaning is a position in this huge semantic space, which is defined relative to all other positions that constitute this space. We thus represent meaning geometrically, i.e. mathematically, which means that we can calculate with meanings. For instance, we can readily calculate how close or far apart two vectors are in this semantic space – hence, the degree of semantic relationship between any words, sentences, or texts.

The technique that allows us to construct such a semantic space is Latent Semantic Analysis (LSA), as developed by Landauer and his coworkers (for introductions, see Landauer, 1998; Landauer & Dumais, 1997; Landauer, Foltz & Laham, 1998). A good way to form an intuition about LSA is to compare it with how people used to make maps (before satellite photographs): they collected a large number of observations about distances between various geographical landmarks and then put all these observations together in a two-dimensional map. Things will not fit perfectly because of measurement errors or missing information, but on the whole, it turns out that we can arrange all the geographical distances in a two-dimensional map, which is very useful because it allows us to calculate distances between points that were never measured directly. Note that if we want to make a map of the world, we will not be able to put all of our data into a two-dimensional map without severe distortions; we need three dimensions for this purpose. LSA constructs semantic spaces in an analogous way. The basic measurements are word co-occurrences. In the case of the semantic space used below, that means over 30,000 documents with over 90,000 different words for a total of about 11 million words. But what should be the dimensionality of the map that is to be constructed? If we employ too few dimensions (two or three, or even 100), the map will be too crude and cannot reflect the kind of semantic relations among words that people are sensitive to. Maps in too many dimensions are not very useful either, however. There is too much accidental, non-essential, even contradictory information in co-occurrence data, because which words are used with other words in any concrete, specific instance will depend on many factors, not just their meaning. We need to discard this excess and focus on the semantic essentials. It turns out, as an empirical fact, that semantic maps – spaces – of 300-400 dimensions yield results that are most closely aligned with human judgments.

LSA thus represents the meaning of a word as a vector in a 300-dimensional semantic space (that is, as a list of 300 numbers that are meaningful only in relation to the other vectors in that space). The meaning of a set of words can be represented as the centroid (vector sum) of the individual word vectors. Thus, sentence meanings are computed as the sum of the words, irrespective of their syntactic structure. Obviously, such a procedure neglects important, meaning-relevant information that is contained in

word order and syntax. In spite of this limitation, LSA has proven to be a powerful and useful tool for many purposes (see the references above). Nevertheless, the neglect of syntax is a serious limitation for LSA which is especially noticeable when we are dealing with short sentences.

The Predication Model of Kintsch (2001) was designed to overcome this limitation, at least for simple argument-predicate sentences. Specifically, the meaning of a predicate is modified to generate a contextually appropriate sense of the word. Consider

The stock market collapsed and *The bridge collapsed.*

The meaning of the predicate *collapsed* that is used here with two different arguments depends on its context: different aspects of *collapse* are foregrounded when the stock market collapses than when a bridge collapses. We say that *collapse* has more than one sense. (There are words, homonyms like *bank*, that have more than one meaning). The Predication Model generates context appropriate senses (or meanings) of a predicate by combining an LSA knowledge base with the construction-integration model of text understanding of Kintsch (1998). It modifies the LSA vector representing the predicate by combining it with features of its semantic neighborhood that are related to the argument of the predication. Specifically, it constructs the semantic neighborhood of the predicate (all the other vectors in the semantic space that are most closely related to the predicate) and then uses a constraint satisfaction process to integrate this neighborhood with the argument: *stock market* selects certain features from the neighborhood of *collapse*, while *bridge* selects different ones. The selected neighborhood vectors are then combined with the predicate vector to yield a context-sensitive sense of the predicate. A more detailed description of this model is given in Kintsch (2001) and the Appendix.

Generating context sensitive word senses does not always produce dramatic results. In the sentence

My lawyer is young

the meaning of *young* is not much modified by *lawyer*. This is different for metaphors. In

My lawyer is a shark

the meaning of the predicate *is-a-shark* is very different from *shark* in isolation – the fishy features of *shark* are de-emphasized (e.g., *has-fins*, *swims*), but they do not

disappear, while other features of shark (e.g., *vicious*, *mean*, *aggressive*) are weighted more strongly because they are somewhat *lawyer*-related, whereas *has-fins* is not.

Kintsch (2000) has shown that this predication algorithm yields interpretations of simple Noun-is-a-Noun metaphors that are in agreement with our intuitions about the meaning of metaphors by comparing the vector generated by the model with appropriate landmarks. The measure used for these comparisons is the cosine of the angle between respective vectors, which can be interpreted in much the same way as correlation coefficients. Thus, the cosine between highly similar vectors is close to +1, while unrelated vectors have a cosine close to 0. For example, *surgeon* is related to *scalpel* (cos=.29) but not to *axe* (cos=.05), while *butcher* is related to *axe* (cos=.37) but not to *scalpel* (cos=.01). *My surgeon is a butcher* moves *surgeon* closer to *axe* (cos=.42) in the semantic space and farther away from *scalpel* (cos=.10). Conversely, *My butcher is a surgeon* relates *butcher* to *scalpel* (cos=.25) and diminishes but does not obviate the relationship to *axe* (cos=.26). Examples like these demonstrate that the LSA space, together with the predication algorithm, represent the meaning of metaphors in a human-like way.

In a recent review, Gibbs (2001) compared several models of figurative language understanding. It is instructive to situate the present approach among current conceptions of metaphor comprehension in psycholinguistics, several of which are closely related to it, while others provide illuminating contrasts. The two models closest to the present approach are the class-inclusion model of Glucksberg (1998) and the underspecification model of Frisson & Pickering (2001). Glucksberg's view that Noun-is-a-Noun metaphors are class inclusion assertions where the appropriate class is newly generated by the metaphor, was the basis for developing the present model in Kintsch (2000). Indeed, LSA and the predication model are one way in which the notion of generating metaphorical superordinate categories can be operationalized. Frisson & Pickering's notion that people initially access an underspecified meaning of words and then elaborate it in context also describes the predication algorithm on which the present model is based. Specifically, the underspecified representation of polysemous words in the present case is the LSA vector (which is not so much *underspecified* as *unspecified*, since it lumps together all meanings and senses of a word); the mechanism that generates a specific,

context appropriate interpretation is the constraint satisfaction process of the predication algorithm. A comparison with the constraint satisfaction model of Katz & Ferretti (2001), on the other hand, points out a limitation of the present model: the spreading activation process (see the example in the Appendix) considers only semantic constraints, while Katz & Ferretti want to consider a broader range of constraints (e.g., syntactic constraints).

Gentner & Bowdles (2001) highlights another limitation of the present approach. Some metaphors are understood like analogies, i.e. by structural alignment, which is a controlled, resource demanding process. The predication algorithm, in contrast, applies when sentences, (metaphorical or not), are understood automatically, without requiring this kind of problem solving.

The principal difference between the present model and other models - psycholinguistic, linguistic, or philosophical - is that it is a fully realized, computational theory. Below we explore whether this computational model arrives at interpretations that are like human interpretations. In Kintsch (2000), the LSA vectors generated by the model were compared with intuitively plausible landmarks. For instance, it was shown that *My lawyer is a shark* is closer to *viciousness* than *lawyer* by itself, which is what one would expect. Here, we employ a method that does not require the use of selected landmarks. Instead, we directly compare the vector constructed by the model with the set of interpretations of a metaphor generated by people. If the model successfully captures the meaning of the metaphor, the sentence vector should be more closely related to the set of interpretations generated by human comprehenders than to the individual words of the sentence.

We also propose to examine the computational processes that generate the vectors for different classes of metaphors for clues as to what differentiates the processing of easy and difficult metaphors. It is well known empirically (Katz et al., 1988) that there are large differences in the ease with which metaphors are understood. What is it that differs when the model processes easy and difficult metaphors? If we observe such a difference, this may be a clue about the sources of comprehension difficulty in human understanding.

Method

Participants

Twenty-four undergraduate students at the University of Colorado participated in the experiment. All were native speakers of English and received class credit for their participation.

Materials and procedure

Each participant was tested individually in a twenty-minute experimental session. After giving informed consent, each participant received an experimental packet consisting of a page of instructions and three pages of stimuli (10 stimulus sentences per page). Each stimulus sentence was a metaphorical statement of the N_1 -is- N_2 (for example, *My lawyer is a shark*.)

Each participant saw the metaphors in the same fixed order. The stimulus order was pseudorandom with the constraints that no two metaphors with the same argument were adjacent and that no more than three easy or three difficult metaphors were presented in a row. The judgment of which metaphors would be easy and which would be difficult was based on data from a pilot experiment using these stimuli.

Beneath each stimulus sentence were two additional items. The first was a sentence completion frame consisting of the subject and verb "X is" of the original metaphor sentence followed by a blank line. Participants were instructed to complete the sentence with a literal version of the original metaphor. For example, if the participant saw the metaphor, *My lawyer is a shark*, followed by *My lawyer is _____* " s/he might fill in *very mean* in order to reflect the literal meaning of the metaphor. After each sentence completion, a set of rating numbers was listed. The participants were asked to circle a number (1-5) to reflect the difficulty of comprehending the stimulus metaphor. A rating of "1" indicated that the metaphor was very easy to understand, and a rating of "5" indicated that the metaphor was very difficult to understand. Participants were instructed to work their way through the packets and to try to come up with an answer and rating for each stimulus metaphor.

Results

Average difficulty ratings were calculated for each stimulus metaphor. Difficulty ratings ranged from 1.29 (*The mosquito is a vampire*) to 4.21 (*A factor is an administrator*). Thirteen metaphors had a rating of 2 or lower and 13 metaphors had a rating of 3 or higher. For the simulations, these were designated as easy and difficult respectively. The remaining four metaphors with intermediate ratings were discarded.

Table 1 shows that the easy and difficult metaphors were clearly differentiated not just in their ratings but also in terms of the interpretive responses subjects generated. For easy metaphors, almost half (48%) of all responses were identical in meaning (e.g. *blood sucker, sucks blood, blood sucking* for the *Mosquito is a vampire* metaphor). Much less agreement (21%) existed among the subjects for difficult items, $t(24) = 4.02, p < .01$. While there were no failures to respond on the generation test for easy items, subjects could not generate a response on 7% of the trials for the difficult items. Furthermore, if one looks at the whole set of responses generated by the subjects, that set was more coherent for easy items than for hard items. The coherence measure used here is the average cosine of each subject's response to the whole set of responses for a particular metaphor, shown in the last column of Table 1. The difference between the coherence of easy items and difficult items was statistically significant, $t(24) = 4.38, p < .01$.

Table 1

The results for difficult metaphors are noteworthy. Faced with items such as *A factor is an administrator* or *Happiness is a ditch*, people don't just give up, but find some interpretation or another. And it is not just a random interpretation, either: on the average about 4 to 5 of our 23 subjects came up with the same response, which is less than the 11-subject agreement we found for easy metaphors, but far from random. Similarly, the responses subjects generated for the difficult metaphors were more diverse (the average cosine between a response and the total response set is .55) than for easy metaphors (average cosine is .64), but what is striking is that there still was a considerable level of agreement, even for what one might regard as pure nonsense.

Results of the Simulations

To determine how well the model was able to fit the data from the rating study, we used the predication model to compute the cosine between the vector representing the meaning of a metaphor and the vector representing the set of all responses generated by the subjects for this metaphor. Averaged over the 26 metaphors used for the simulation, this cosine has a value of .51. There is, however, no way of determining whether this value is high or low, for the absolute value of cosines in LSA depends on many factors; only relative values for cosines computed in the same way can be readily compared. Table 2 provides such a comparison. The cosine between the metaphor vector computed by the model and the set of responses generated by the subjects is higher than the cosine between either:

- a) just the argument of the metaphor and the set of subject responses OR
- b) just the predicate of the metaphor and the set of subject responses,

$p < .001$ by sign test. Importantly, the cosine generated by the predication model is also higher than the cosine between the centroids of the argument and predicate of a metaphor, $p < .001$ by sign test. Thus, while we cannot claim that the model predictions are good on an absolute scale, we know that they are better than what can be achieved by either the predicate or argument alone, or by the centroid of the two.

Table 2 also shows that there was no difference between how well the model fits the subjects' responses for easy and hard metaphors. It is clearly not the case that the model fits the data only when subjects agree with each other, that is, for the easy metaphors. When people agree about the interpretation of a metaphor the model computes a vector that is closely related to that agreed upon meaning of the metaphor. However, for difficult metaphors, where there is much less agreement and subjects generate a more diverse set of responses, the vector computed by the model is just as close to the average of the subjects' responses. For easy metaphors, the model focuses in on some specific meaning; for difficult metaphors, it specifies a diffuse but non-arbitrary meaning – just as real people do. To understand what is happening here one must remember that the LSA vector for a set of a responses is the centroid, i.e. average, of the individual response vectors. The model vector is equally close to that average of easy

and the average of difficult items, although the average for the easy items is computed from a narrow range of responses while the average of the difficult items is based on a diffuse set of responses.

Table 2

This interpretation is supported by an analysis of the relationship between the metaphor vectors computed by the model and the modal responses given by the participants. As Table 1 shows, almost half of all responses were common for easy metaphors; the average cosine between these modal responses and the metaphor vector is .32ⁱ. In contrast, many fewer common responses were given to difficult metaphors (21%), and their cosine with the metaphor vector is significantly lower, cosine = .22 ($t(24) = 1.75$, $p < .05$). The model calculates a vector that is equally close in semantic space to the set of responses participants produce for easy and difficult metaphors. However, that set is different for easy and difficult metaphors. For easy metaphors, there is agreement among participants and their choice is strongly related to the model vector. For difficult metaphors, there is much less agreement among participants, and their choice is less strongly related to the model vector. In the first case, the vector is fairly precise and generally focuses on a particular concept – the modal response of the participants; in the second case, this focus is lacking, but the model vector nevertheless captures the variety of responses produced by the participants.

Given these results, the question arises, “what makes a metaphor easy or difficulty according to the predication model?” One obvious candidate is the semantic distance between the argument and predicate of a metaphor. One might suppose that if the two terms are very far apart in the semantic space, it might be difficult to find something they have in common. This conjecture does not hold up, however: the average cosine between the argument and predicate for easy metaphors is .10, versus .07 for difficult metaphors – a difference that is unreliable statistically, $t(24) = .96$. Thus, metaphors are not difficult because their argument and predicate terms are unrelated overall.

Another possibility is that processing difficulty depends on how much information is available about either of the two terms of a metaphor. However, the data do not support this hypothesis either. The length of a vector is a measure of how much information LSA has about a word. The average vector length for the predicates of the

easy metaphors was .86, versus 1.27 for the hard metaphors, $t(24) = 1.12$, $p > .05$. Another way of measuring how much LSA knows about a word is to look at the number of other words that are close neighbors. However, there is no difference between the predicates of easy and hard metaphors in this respect either: easy predicates have on the average 17 neighbors with a cosine greater than .5 and 36 neighbors with a cosine greater than .4, and hard predicates have 18 neighbors with a cosine greater than .5 and 22 with a cosine greater than .4. Finally, there is no difference in the vector length of the arguments of easy (1.19) and hard metaphors (1.00), $t(24) = .63$, $p > .05$. Thus, it does not appear that processing difficulty is related to properties of either the argument or the predicate of a metaphor in isolation.

Table 3

A more promising hypothesis is that processing difficulty depends on whether at least a few items can be found that are strongly related to both the argument and the predicate of a metaphor. As described in the Appendix, the predication model for metaphors of the form N_1 -is- N_2 works by selecting neighbors of N_2 that are most closely related to N_1 and uses these terms to modify the N_2 -vector. Are these terms more closely related to N_1 and N_2 for easy metaphors than for difficult metaphors? It turns out that is not the case for N_2 – the average cosine between N_2 and the set of selected terms is not significantly different for easy and hard metaphors, $t(24) = .62$, as shown in Table 3. However, the differences between easy and difficult metaphors is statistically significant for N_1 , $t(24) = 2.22$, $p \leq .05$, and marginally significant for $N_1 * N_2$, $t(24) = 1.54$, $p = .07$. Thus, the predication model suggests that metaphors are easy to process if the argument has a good match among the close neighbors of the predicate; if the match is less good, this is experienced as processing difficulty, perhaps because the search for a better match continues into regions where items are no longer sufficiently strongly related to N_2 . This must remain a tentative conclusion, however, for the relationship is not overly strong: the correlation between the rated difficulty of a metaphor and the cosine(selections : N_1) is only $r = -.46$, which is significant statistically, but not very high.

Conclusions

Latent Semantic Analysis or LSA allows one to represent the meaning of words as vectors in a high-dimensional semantic space. The meaning of sentences can be computed from these vectors. There are two ways of computing a sentence vector, given the constituent word vectors. One is context free. It disregards the syntax and simply sums up the word vectors. Another possibility is to adjust the word vectors contextually according to their syntax. Specifically, in the predication model of Kintsch (2001), sentence vectors of the form N_1 -is- N_2 are computed by modifying the predicate vector N_2 according to the argument vector N_1 . Thus, a context appropriate sense of the predicate is generated. This model has been shown to provide a good account of several semantic phenomena that otherwise are outside the scope of LSA by itself, including metaphor comprehension (Kintsch, 2000). The present results provide further evidence that the predication model is capable of adequately representing the meaning of simple nominal metaphors, in the sense that the metaphor vectors it computes are closely related to the interpretations that people give to these metaphors.

The data reported here focus on differences in the way people interpret easy and difficult metaphors. We have shown that metaphors that are judged to be easy to comprehend are interpreted in similar ways by most people, whereas a greater range of interpretations exists for difficult to comprehend metaphors. However, people agreed among each other to some extent, even when the metaphors they were asked to interpret appeared to be pure nonsense. Faced with the seemingly impossible task of finding an interpretation for such metaphors, people do not give up (a failure to respond was observed in only in 7% of the cases); instead they come up with something, and there is a certain consistency among people in how they respond. There is not nearly as much consistency for difficult metaphors as for easy metaphors, but while interpretations are diffuse and vague for difficult metaphors, they are not random. This consistency in people's responses may, however, not derive from having successfully interpreted a difficult metaphor, but may simply reflect word-based constraints.

Interestingly, the predication model behaved in much the same way: it came up with vague and less coherent interpretations for difficult metaphors, but it matched what

people said as well as for easily comprehended metaphors. For easy metaphors there is widespread agreement among people, and the model produces a vector close to that agreed upon interpretation. For difficult metaphors, responses are more varied, but the model produces a vector that is just as close to these varied responses as it is to the generally agreed upon interpretation of a good metaphor. For both people and model there is something in the semantic structure that guides their interpretation. The semantic structure provides a tight constraint for easy metaphors, and only a loose one for hard metaphors, but the comprehension process neither collapses nor becomes random.

If the model understands easy and difficult metaphor equally well (in the sense that it predicts human interpretations equally well), then what is different about the computational process for easy and difficult metaphors? It is not the case that the constituent words by themselves are more or less informative, nor is it the case that easy understanding requires a pre-existing global relation between the two terms of a metaphor. Rather, it appears that, although argument and predicate can be totally unrelated overall, a metaphor is comprehensible if some link is found between topic and vehicle, even though the two may be unrelated overall. Thus, *lawyer* and *shark* are orthogonal in the semantic space ($\text{cosine}(\textit{shark}:\textit{lawyer}) = -.01$), but there are certain aspects – like *vicious* or *mean* – that link the two and make “*My lawyer is a shark*” an easily understandable metaphor.

Theories of metaphor comprehension have traditionally been informal. We hope that by offering a formal model that can yield quantitative experimental predictions and, at the same time, is conceptually related to the issues under discussion in the psycholinguistic literature, further progress can be made in our understanding of comprehension processes, metaphoric as well as literal. We also claim that the results presented here show that LSA provides a useful basis for a psychological theory of meaning.

References

- Cacciari, C. & Glucksberg, S. (1994). Understanding figurative language. In M. A. Gernsbacher (Ed.), *Handbook of psycholinguistics* (pp.447-478). San Diego: Academic Press.
- Frisson, S. & Pickering, M. J. (2001). Obtaining figurative interpretations of a word: Support for underspecification. *Metaphor and Symbol, 16*, 133-187.
- Gentner, D. & Bowdle, B. F. (2001). Convention, form, and figurative language processing. *Metaphor and Symbol, 16*, 223-247.
- Gibbs, R. W. Jr. (1994) Figurative thought and figurative language. In M. A. Gernsbacher (Ed.), *Handbook of psycholinguistics* (pp.441-446). San Diego: Academic Press.
- Gibbs, R. W. Jr. (2001). Evaluating contemporary models of figurative language understanding. *Metaphor and Symbol, 16*, 317-333.
- Glucksberg, S. (1998). Understanding metaphors. *Current Directions in Psychological Science, 7*, 39-43.
- Katz, A. N. & Ferretti, T. R. (2001). Moment –by-moment readings of proverbs in literal and nonliteral contexts. *Metaphor and Symbol, 16*, 193-221.
- Katz, A. N., Paivio, A., Marschark, M., & Clark, J. M. (1988). Norms for 204 literary and 260 nonliterary metaphors on 10 psychological dimensions. *Metaphor and Symbolic Activity, 3*, 191-214.
- Kintsch, W. (1998). *Comprehension: A paradigm for cognition*. New York: Cambridge University Press.
- Kintsch, W. (2000). Metaphor comprehension: A computational theory. *Psychonomic Bulletin & Review, 7*, 257-266.
- Kintsch, W. (2001). Predication. *Cognitive Science, 25*, 173-202.
- Landauer, T. K. (1998). Learning and representing verbal meaning: Latent Semantic Analysis theory. *Current Directions in Psychological Science, 7*, 161-164.
- Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The Latent Semantic Analysis theory of acquisition, induction and representation of knowledge. *Psychological Review, 104*, 211-240.

Landauer, T. K., Foltz, P., & Laham, D. (1998). An introduction to Latent Semantic Analysis. *Discourse Processes*, 25, 259-284.

Table 1

Rated difficulty and properties of the responses generated for easy and difficult metaphors

	Rated Difficulty of Comprehension	Modal Response Frequency	No-Response Frequency	Coherence of Responses
Easy	1.75	48%	0%	.64
Difficult	3.68	21%	7%	.55

Table 2

Cosines between the vectors for metaphors, their argument (N_1) and predicates (N_2), and all the responses generated by subjects

	Cosine{metaphor : responses}	Cosine{ N_1 : responses}	Cosine{ N_2 : responses}
Easy	.50	.34	.34
Difficult	.51	.34	.31

Table 3

Average cosines between the items selected by the model to modify the predicate-vector and the predicate (N_2) and argument (N_1) of the metaphor, as well as the product of the cosines

	$\cos\{\text{selections: } N_2\}$	$\cos\{\text{selections: } N_1\}$	$\cos\{\text{sel: } N_1\} * \cos\{\text{sel: } N_2\}$
Easy	.40	.30	.12
Difficult	.43	.21	.08

Acknowledgements:

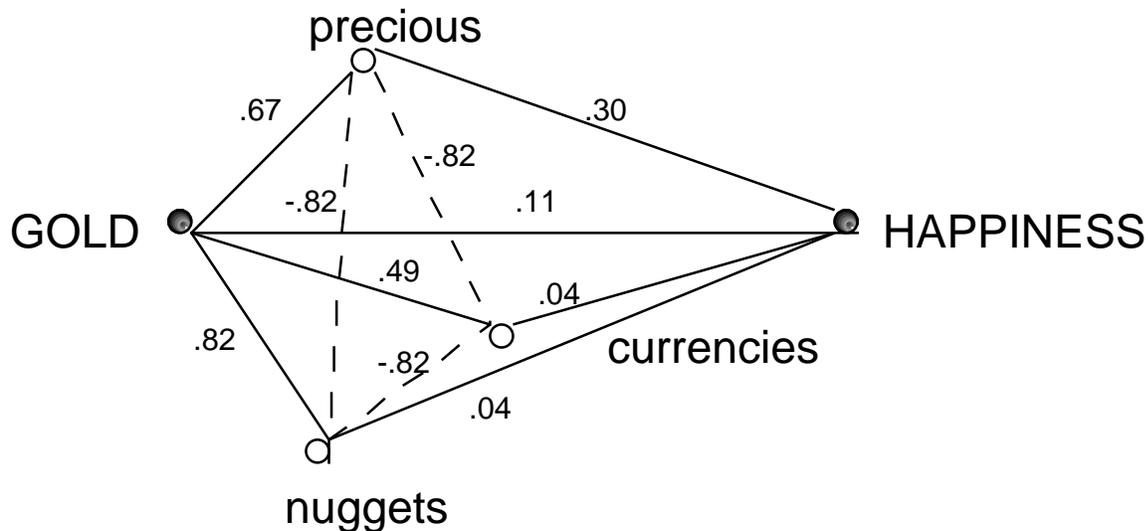
This research was supported by a contract from the Army Research Institute to T. K. Landauer and W. Kintsch. We thank Tom Landauer, Kirsten Butcher, Eileen Kintsch, and the LSA Research Group for their help.

Appendix

The Predication Algorithm (after Kintsch, 2001)

In sentences of the form Noun₁-is-Noun₂, Noun₁ is called the argument (A) and Noun₂ is called the predicate (P); the word *is* is neglected. The meaning of A and P is represented in LSA as the vectors **A** and **P**. Normally, in LSA the meaning of the sentence A-is-P is given by the vector sum **A** + **P**. According to the predication algorithm the meaning of the sentence A-is-a-P is given by the vector sum **A** + **P_A**, where **P_A** is the contextually modified predicate vector. To calculate **P_A** the construction-integration (CI) model of discourse comprehension of Kintsch (1998) is used to select from the semantic neighborhood of P those items that are in some way relevant to A. This selection is achieved through a spreading activation process: a network is constructed consisting of P and A and the closest neighbors of P. Activation is spread in that network. The most strongly activated neighbors of P will be used to modify P to create **P_A**.

In the calculations reported here, A was allowed to select those items most relevant to it from the 500 closest neighbors of P. However, the details of this process are best described with a more manageable example. Consider the metaphor *Happiness is gold*. The following network consisting of P (*gold*), A (*happiness*), and three close neighbors of P (*precious*, *currencies*, and *nuggets*) can be constructed:



The positive link strengths are the cosines between the respective items in the LSA space. Thus, for example, *happiness* and *gold* have a $\cos = .11$, that is they are only weakly related. The neighbors of P interfere with each other, that is, they compete for the activation in the network. Interference is generated by linking the three neighbors with negative links, so that the sum of the positive links equals the sum of the negative links in the network. Under these conditions, the neighbors of P with the highest cosines with A will become activated in the net, while neighbors that are only related to P but not to A will become deactivated. In the present example, the activation of the five nodes in the network when the process stabilizes areⁱⁱ

<i>Gold</i>	1.000
<i>Precious</i>	1.000
<i>Currencies</i>	0.000
<i>Nuggets</i>	0.048
<i>Happiness</i>	1.000

Thus, *currencies* and *nuggets* (in spite of its very strong relationship to *gold*) have been rejected and *precious* has been selected to modify **P**. Let **N** be the LSA vector for *precious*. Then $\mathbf{P}_A = \mathbf{P} + \mathbf{N}$.

For the calculations reported here, a network with the 500 closest neighbors of P is needed. Because of the large size of this network, an approximation procedure described

in Kintsch (2001) was used. The outcome of the spreading activation process was estimated by rank ordering the product of each item's cosine with P and its cosine with A and selecting the five largest products to calculate \mathbf{P}_A . This estimation procedure and the spreading activation algorithm usually, but not necessarily, select the same set of neighbors. In the example above, the link products for *precious*, *nuggets*, and *currencies* are .20, .04, and .03, respectively; i.e., the item with the highest value is *precious*, which was also the most highly activated neighbor in the original spreading activation process.

ⁱ The absolute value of the cosines between the metaphor vector and the total response set on the one hand and the modal response on the other are not comparable, since the former involves a comparison between two sets of words whereas the latter compares a single word with a set of words.

ⁱⁱ For this illustration, the activation of *gold* and *happiness* was fixed at a value of 1. In large networks these activation values are allowed to fluctuate. When a network settles in a state where either of these nodes is not among the 5 most highly activated ones, this indicates a failure of the comprehension process (i.e., the algorithm failed to find links between P and A.)