



Published in final edited form as:

J Mem Lang. 2001 October 1; 45(3): 433–450. doi:10.1006/jmla.2000.2782.

The Role of Suppression and Enhancement in Understanding Metaphors

Morton Ann Gernsbacher,
University of Wisconsin-Madison

Boaz Keysar,
University of Chicago

Rachel R. W. Robertson, and
University of Wisconsin-Madison

Necia K. Werner
University of Wisconsin-Madison

Abstract

Participants read either a metaphorical prime sentence, such as *That defense lawyer is a shark*, or they read a baseline-prime sentence. The baseline-prime sentence was literally meaningful in Experiment 1 (e.g., *That large hammerhead is a shark*), nonsensical in Experiment 2 (e.g., *His English notebook is a shark*), and unrelated in Experiment 3 (e.g., *That new student is a clown*). After reading the prime sentence, participants verified a target property statement. Verification latencies for property statements relevant to the superordinate category (e.g., *Sharks are tenacious*) were faster after participants read the metaphor-prime sentence than after they read the baseline-prime sentence, producing an enhancement effect. In contrast, verification latencies for property statements relevant to only the basic-level meaning of the vehicle and not the superordinate (e.g., *Sharks are good swimmers*), were slower following the metaphor-versus the baseline-prime sentence, producing a suppression effect. As Glucksberg and Keysar's (1990) class inclusion theory of metaphor predicts, the enhancement effect demonstrates that the vehicle of a metaphor stands for the superordinate category of the vehicle, and the suppression effect demonstrates that the metaphorical vehicle does not stand for its basic-level meaning.

Language users experience little difficulty understanding a sentence such as *Perjury is a boomerang*, especially if they followed trials involving a perjured president. In contrast, theories of language have particular difficulty when it comes to explaining how language users understand metaphorical sentences. This difficulty stems from the traditional literal bias of theories of language.

A literal interpretation of *perjury is a boomerang* would yield an anomaly, either on semantic or pragmatic grounds; perjury is not a boomerang, and attempting to categorize *perjury* into *boomerang* as one categorizes robins as birds results in a semantic anomaly or

deviance (e.g., Altwerger & Strauss, 1987; Davidson, 1978; Grice, 1975; Kittay, 1987; Levin, 1977; Searle, 1993). To date, though, psycholinguistic investigations have revealed no evidence that comprehenders experience such an anomaly or so called semantic deviance when comprehending metaphors (e.g., Gerrig, 1989; Gibbs, 1984; Glucksberg, Gildea & Bookin, 1982; Keysar, 1989, Keysar, 1994, Ortony, Schallert, Reynolds & Antos, 1978). Why not? Why don't people experience such a deviance, even as an intermediary product of the comprehension process?

The answer depends on one's theory of metaphor understanding. Some theories argue that people do not consider these statements to be categorical at all. If the sentence does not express a category relation, then there is no categorical mistake and, therefore, no anomaly is involved (e.g., Ortony, 1979). Alternative theories (Glucksberg, 1998; Glucksberg & Keysar, 1990; Glucksberg & McGlone, 1999) propose that such statements are indeed understood as category assertions. They propose that people interpret the metaphor by categorizing *perjury* into *boomerang*. However, according to Glucksberg and Keysar (1990) the term *boomerang* does not refer to the object boomerang; instead, it refers to a higher level category that could include perjury. The experiments reported in this paper tested this aspect of the theory.

Approaches to the Study of Metaphor

The cognitive study of metaphor focuses on a variety of issues, not always explicitly distinct from each other. Analogical problem solving (Bassok, in press; Carbonell, 1986; Clement & Gentner, 1991; Gick & Holyoak, 1983; Holyoak & Thagard, 1989; Sternberg, 1977; Thagard, 1988) may or may not involve the same processes as comprehending metaphors. Understanding how people store and remember metaphors is a different pursuit than understanding how people comprehend them (Marschark, Katz & Paivio, 1983; Marschark & Hunt, 1985). What makes a metaphor good or apt (Malgady & Johnson, 1976; Tourangeau & Sternberg, 1981) is a different question from what causes people to categorize a sentence as metaphorical and how they judge its degree of metaphoricality. Similarly, as Gerrig and Healy (1983) argued, metaphor understanding differs from metaphor evaluation. Our focus in this paper is on the processes that underlie the understanding of mundane nominative metaphorical assertions (e.g., *That defense lawyer is a shark*). Our question, simply put, is the following: What do people do to make sense of a metaphorical sentence?

Different theories make different assumptions regarding the processes that underlie metaphor understanding. One major approach, which goes back to the historical "comparison" theory of metaphor, assumes that an act of comparison is at the heart of the process. This approach has a variety of instantiations. For instance, Miller's theory (1993) assumes that understanding the metaphor, *man is a wolf* requires translating it to the simile, *man is like a wolf*. The assumption is that the metaphor conceals an implied comparison that has to be uncovered as part of the understanding process. Miller proposes a series of rules that allow, by analogy, the comparison of *man* to *wolf* to yield an interpretation.

Comparison is also at the heart of Ortony's (1979) theory of salience imbalance. According to this theory, understanding *cigarettes are time bombs* involves comparing the topic, *cigarettes*, to the vehicle, *time bombs*, to reveal the highly salient properties of *time bombs* that are low in salience in *cigarettes*; these differentially salient properties are the ground for the metaphor. In Gentner's theory (e.g., Gentner, Falkenhainer & Skorstand, 1988; Wolff & Gentner, 1992), comparison is important in the early matching process of the topic and the vehicle, a matching process that looks for features and relations common to the topic and vehicle. Kintsch (2000) has further drawn upon these comparison theories by representing the interaction between the meaning of the topic and the vehicle in a computational model of metaphor comprehension.

Tourangeau and Sternberg's (1982) domain interaction theory assumes a different kind of comparison — a comparison not between the topic and the vehicle terms, but instead between the two domains from which they come. According to Tourangeau and Sternberg, understanding the metaphor *the eagle is a lion* involves comparing the domain of mammals to the domain of birds to discern the relative position of *lion* among *mammals* and applying that relation to *eagle* in its domain of *birds*.

As an alternative to comparison-based theories, Glucksberg and Keysar (1990) propose a categorization-based approach, which does not assume a comparison between terms. Next, we shall describe the basic assumptions of Glucksberg and Keysar's (1990) theory, focusing on the specific predictions that the experiments reported here tested.

The Class Inclusion Theory of Metaphor

Metaphors and ad-hoc categories—Glucksberg and Keysar (1990; see also Glucksberg, 1991; 1998; Glucksberg & Keysar, 1993; Glucksberg, Manfredi, & McGlone, 1997; McGlone & Manfredi, in press) proposed that when people comprehend metaphors, such as *rage is a volcano*, they do not compare *rage* and *volcano*. Instead, they understand the sentence as a categorization statement. According to the theory, to understand metaphors, people construct an ad-hoc category (Barsalou, 1983) that the vehicle typifies. In this case, *volcano* typifies the “category of things that erupt unexpectedly and might cause damage.” This category includes a variety of members such as epidemics, revolutions, rage, and so on. The metaphor, then, is understood as an assertion that *rage* is a member of this category.

If the vehicle is the basis for constructing the ad-hoc category, one might conclude that the exact same category will be constructed for all metaphors that contain the same vehicle, regardless of the identity of the topic. This is not the case; the topic affects the nature of the category. For example, consider the difference between *my dentist is a magician* and *my stock broker is a magician*. Even though both metaphors have *magician* as their vehicle, we might mean in one case that the dentist is able to make an expected event disappear (e.g., pain), whereas in the other case the broker is able to create something (e.g., profit) out of nothing.

Glucksberg, McGlone, and Manfredi (1997) suggest that the topic constrains the nature of the constructed category because it provides possible dimensions for attribution. For

example, the *dentist*-topic might provide relevant dimensions for attribution such as “degree of manual skill,” “ability to control level of pain,” and so on. The *stock broker*-topic provides dimensions such as “ability to make accurate predictions,” “timing ability,” and so on. In contrast, the vehicle of the category, by creating the relevant category, provides a value on the suggested dimension. Glucksberg, McGlone, and Manfredi (1997) tested these ideas and demonstrated that the topic and the vehicle play two different roles in the course of metaphor understanding. Our focus in this paper is on the role of the vehicle in the understanding process.

The vehicle stands for the superordinate category—The strong version of the class inclusion theory assumes that not only does the vehicle typify the created category, but it lends its name to it. This possibility was briefly suggested by Roger Brown as well (1958):

“metaphor differs from other superordinate-subordinate relations in that the superordinate is not given a name of its own. Instead, the name of one subordinate (i.e., the vehicle) is extended to the other” p. 140.

Such a naming device, using a typical member to name the superordinate category, is not uncommon in language. There are many cases where language users name a category that has no conventional name by using its typical member or members. For example, in American Sign Language one can refer to the superordinate *furniture* by signing “chair-table-bed, etc.,” all of which are typical furniture (Newport & Bellugi, 1978). The use of a prototypical member to name the category is also common in American Indian Languages. For example, the name for eagle is used by Shoshoni speakers to refer to the category of large birds (Hage & Miller, 1976; for other examples, see Glucksberg & Keysar, 1990). So, according to the class inclusion theory of metaphor, when we say that *her rage was a volcano* we mean that her rage was a member of that category that has the name *volcano*. *Volcano* in this context no longer stands for a physical object, but it stands for the entire category that it typifies.

If we consider *volcano* as in *Mount St. Helen is a volcano* to be in the basic-level of categorization, then *volcano* as a metaphor vehicle is in the superordinate-level of categorization. The strong version of the class inclusion theory assumes that when we understand the metaphor *rage is a volcano*, we do the same kind of thing that we do when we understand the assertion *a robin is a bird*. We categorize either *rage* or *robin* into their corresponding superordinate category, *volcano* and *bird*. The only difference is that *bird* is a conventional name for a taxonomic category, whereas *volcano* is a borrowed name for an ad-hoc category that is created on the fly.

This category-based approach can account for a variety of metaphor phenomena (Glucksberg & Keysar, 1990). For instance, metaphor is usually irreversible because reversing the terms typically yields an anomaly (Ortony, 1979). It doesn't make sense to reverse the metaphor *perjury is a boomerang* to **a boomerang is perjury*. A class inclusion theory naturally accounts for this irreversibility: Irreversibility is inherent in any class inclusion statement; it does not make sense to reverse *a robin is a bird* to *a bird is a robin*. Because the theory conceives of *rage is a volcano* as a category statement, it would not make sense to reverse it for the same reason. Indeed, Glucksberg, McGlone and Manfredi

(1997) demonstrated that people perceive metaphors not merely as asymmetrical but as irreversible.

The class inclusion theory can even predict when a reversed metaphor would yield a meaningful statement. If the topic typifies a potential ad-hoc category into which the vehicle can be categorized, then the reversed metaphor would make sense. For example, one can reverse the metaphor *this lion is a shark* to a still meaningful metaphor, *this shark is a lion*. The reversal is not an anomaly because *lion* happens to typify a category of royal creatures. Indeed, with *shark* as the vehicle, the metaphor attributes to the topic a different quality (tenacity) than with *lion* as a vehicle (royalty).

The class inclusion theory of metaphor understanding can also explain why metaphorical comparison statements can be converted into category statements but literal comparisons cannot. It doesn't make sense to convert the literal comparison *cigarettes are like cigars* to the category statement **cigarettes are cigars*. In contrast, it makes a lot of sense to convert the metaphorical comparison *cigarettes are like time bombs* into the category statement *cigarettes are time bombs*. This ability to convert a simile into a metaphor falls naturally out of the theory. A simile can be expressed in a class inclusion form because the simile is an implied category statement (Glucksberg & Keysar, 1990).

The experiments we report here were designed to test the assumption of the class inclusion theory regarding the status of the vehicle. Specifically, we tested the assumption that the vehicle of a metaphor is understood as the superordinate of an ad-hoc category. So *time bombs* in *cigarettes are time bombs* does not stand for the basic-level, the object time bombs. Instead, it names a superordinate ad-hoc category that includes things such as cigarettes and time bombs. We empirically tested this assumption by modifying a paradigm developed by Gernsbacher and her colleagues that identifies two general cognitive mechanisms involved in language comprehension, namely the mechanisms of enhancement and suppression.

Comprehension Mechanisms and Metaphor Understanding

The role of enhancement and suppression in comprehension—Comprehension is enabled by modulating activation; two mechanisms provide this modulation: enhancement and suppression. Enhancement is the increase in activation of memory nodes that represent information central to the on-going comprehension. Suppression is the active reduction in activation of activated memory nodes that represent information that is potentially confusing or irrelevant for comprehension. Gernsbacher and her colleagues have demonstrated the role of these two mechanisms in lexical access (e.g., Gernsbacher & Faust, 1991b; Gernsbacher & St. John, in press), anaphora (e.g., Gernsbacher, 1989; 1997a), cataphora (e.g., Gernsbacher & Jescheniak, 1995; Gernsbacher & Shroyer, 1989), as well as the comprehension of non-linguistic information (e.g., Gernsbacher & Faust, 1991a).

Consider, for example, the role of enhancement and suppression in modulating the access of word meaning. Often the understanding of polysemous words initially involves the activation of contextually relevant as well as irrelevant meanings; however, after a short latency only the relevant meaning remains activated (e.g., Swinney, 1979). According to

Gernsbacher, the mechanism of suppression is what eliminates contextually irrelevant meanings. Gernsbacher and Faust (1991b) demonstrated that the reduction in activation of inappropriate meanings is not due to decay over time, or to compensatory inhibition, but instead to active suppression. Moreover, Gernsbacher and her colleagues discovered that a major element in comprehension skill differences is in the operation of the suppression mechanism (Gernsbacher, 1997b; Gernsbacher & Faust, 1991a; 1994; Gernsbacher, Varner & Faust, 1990; Gernsbacher & Robertson, 1995): While less-skilled comprehenders are as able as more-skilled comprehenders to use contextual information to enhance relevant information, they are less able to suppress irrelevant information.

Enhancement and suppression in metaphor understanding—In the experiments reported here, we investigated the role of the mechanisms of suppression and enhancement in metaphor understanding. Consider first an experiment with polysemous words reported in Gernsbacher (1994). In this experiment participants read sentences and were required to decide quickly if the sentences made sense. Occasionally, consecutive pairs of sentences were designed as prime and target sentences. For example, one target sentence was *She blew out the match*. This target sentence was preceded by one of three prime sentences: a prime sentence using the word *match* with (potentially) the same meaning as the target, *She saw the match*; a prime sentence using the word *match* with a clearly different meaning, *She won the match*; or a nonsense-prime sentence that also contained the word *match*, *She prosecuted the match*.

Participants more rapidly responded that the target sentence made sense after they read the same-meaning prime sentence than after they read the nonsense-prime sentence. If we consider the nonsense-prime sentence as a baseline, these data suggest that the contextually appropriate meaning of a polysemous word is enhanced by reading the word in a biasing context. In contrast, participants more *slowly* responded that the target sentence made sense after they read the different-meaning prime sentence than after they read the nonsense-prime sentence. These data suggest that the other meanings of a polysemous word are actively suppressed by reading the word in a biasing context. The design and rationale of our experiments were analogous to this study.

Our goal was to test the following premise of the class inclusion theory of metaphor understanding: In a metaphor such as *that defense lawyer is a shark*, the vehicle stands for the superordinate category. Therefore, when understanding this metaphor, the vehicle, *shark*, is not interpreted as referring to a basic-level of categorization; rather the vehicle, *shark*, is interpreted as a superordinate ad-hoc category of tenacious, perhaps vicious, things. We tested this proposal by examining whether the superordinate meaning of a vehicle is enhanced once a metaphor is understood. Participants read either a metaphorical prime sentence, such as *That defense lawyer is a shark*, or they read a baseline prime sentence, such as *That large hammerhead is a shark*, where “shark” was used literally. After reading either the metaphor-prime sentence or the literal-prime sentence, participants responded to a target property statement. The property was relevant to the superordinate category, for example, *Sharks are tenacious*. If the metaphorical vehicle stands for the superordinate category, then we should have observed an enhancement effect: Participants should have

more rapidly verified the superordinate property statement after they read the metaphor-prime sentence than after they read the literal-prime sentence.

While such an enhancement effect is a necessary condition for our argument, it is not sufficient to conclude that the metaphorical vehicle stands for a superordinate category. By itself it would not tell us if the vehicle is interpreted at the basic-level or superordinate-level of categorization. The reason an enhancement effect would not be diagnostic is that the (ad-hoc) superordinate meaning of say, *shark*, is inherent in its basic-level of categorization. Being tenacious, which is a central element of the superordinate, ad-hoc category named *shark*, is also central for the basic-level meaning of *shark*. This is not surprising, because the reason *shark* can stand for the superordinate category is that it typifies tenacity. So, while an enhancement effect would not be diagnostic, it is necessary for our claim. If we don't observe such an effect, this would be a serious problem for the class inclusion theory, or for virtually any theory of metaphor.

In contrast to an enhancement effect, a suppression effect would be diagnostic for discriminating between the two levels of categorization. According to the class inclusion theory of metaphor, when people understand the metaphor *that defense lawyer is a shark* they construct an ad-hoc category that the basic-level meaning of *shark* typifies. This means that the process of metaphor understanding makes use of the basic-level meaning of *shark*, and, therefore, that concept should be activated in memory. Indeed, Blasko and Connine (1993) provided evidence for the activation of the literal meaning of metaphor vehicles as people hear metaphorical sentences. Using a cross-modal priming paradigm, they demonstrated that immediately following the vehicle of the metaphor, *hard work is a ladder*, the concept *rungs* was activated. Blasko and Connine (1993)'s finding suggests that the basic-level meaning of *ladder* was available for the readers during comprehension.

However, according to the class inclusion theory, metaphorical vehicles do not stand for literal or basic-level meanings; they represent superordinate concepts. In this sense, a metaphorical vehicle, such as *shark*, is akin to a polysemous word. And just as the contextually inappropriate meaning of a polysemous word is suppressed after the word is disambiguated, properties of the basic-level meaning should be suppressed after a metaphorical vehicle is understood. We therefore predicted that understanding a metaphor involves suppressing the basic-level meaning of the vehicle.

We tested this prediction in the following way: We presented the same two prime sentences as we presented to test our prediction about the role of enhancement in metaphor understanding (i.e., a metaphor-prime sentence, such as *that defense lawyer is a shark*, or a literal-prime sentence, such as *that large hammerhead is a shark*). After reading either the metaphor-prime sentence or the literal-prime sentence, participants responded to a target property statement. In contrast to the property statements that we presented to test our prediction about the role of enhancement in metaphor understanding (i.e., property statements of the superordinate ad-hoc category, e.g., *Sharks are tenacious*), the property statements that we presented to test our prediction about the role of suppression in metaphor understanding were relevant to the basic-level category but not the superordinate ad-hoc category. An example is the property statement, *Sharks are good swimmers*. While being a

good swimmer is a property of the basic-level meaning of *shark*, it is not relevant to the superordinate, ad-hoc category named *shark*.

One of three outcomes was possible. Participants might have read the basic-level property statements *faster* following the metaphor-prime sentences than following the baseline prime sentences. Such enhancement would suggest that the vehicle in the metaphor is not interpreted at the superordinate ad-hoc category level, but instead that it stands for the basic-level *shark*. Such an outcome would be counter-evidence for a class inclusion theory of metaphor. A second possibility was that there would have been no difference in participants' verification latencies to the basic-level property statements following the two types of prime sentences. One interpretation of such an outcome would be consistent with a weaker version of the class inclusion theory, namely, that the vehicle plays a dual role, both at the basic-level and the superordinate-level of categorization.

A third possibility would have supported a strong version of the class inclusion theory: If participants responded *more slowly* to the basic-level property statements after reading a metaphor than after reading the literal (baseline) prime sentence, this would demonstrate a suppression effect. This suppression effect would imply that the basic-level meaning of *shark* was actively suppressed and that the vehicle was interpreted to stand for the superordinate category.

EXPERIMENT 1

Methods

Materials—Participants read a series of sentences. The participants' task was to read each sentence and decide whether it made sense. Half the sentences did make sense, such as *That large hammerhead is a shark*, and half the sentences did not make sense, such as *That apple was a tennis player*. Unknown to the participants, the list of 384 sentences included 48 pairs of experimental sentences (25% of the sentences). We called these sentences “experimental pairs” because the first sentence of each pair was our prime sentence, and the second was our target sentence. We manipulated the prime sentences and measured the effect of our manipulation on participants' responses to the target sentences.

The prime sentences were all in the form of *X is a Y*. Half were metaphorical, such as *That defense lawyer is a shark*, and the other half were literal. For the experimental prime sentences, we first constructed the metaphorical prime sentences and then constructed their literal-prime sentence counterparts by selecting a member of the basic-level category (e.g., *hammerhead*) represented by the metaphorical vehicle (e.g., *shark*) and substituting it for the topic of the corresponding metaphor-prime sentence. For example, for the metaphor-prime sentence, *That defense lawyer is a shark*, the corresponding literal-prime sentence was *That large hammerhead is a shark*.

Each prime sentence was followed by a target sentence. The target sentences were all property statements, such as *Sharks are tenacious*; *Airplanes have wings*; *Vampires suck blood*; *Earthquakes shake buildings*. For the experimental target sentences, the property was relevant to the superordinate category represented by the vehicle of the preceding metaphor

prime, or the property was relevant to only the basic-level category represented by the vehicle of the preceding metaphor prime sentence. For example, for the metaphor prime sentence, *That defense lawyer is a shark*, the property statement relevant to the superordinate category was *Sharks are tenacious*, and the property statement relevant to only the basic-level was *Sharks are good swimmers*. The property statements relevant to the superordinate category were written to highlight what we judged to be the most salient feature of the corresponding metaphor prime sentence. The property statements relevant to the basic-level category were written to highlight what we judged to be the most salient feature of the corresponding literal prime sentence, avoiding statements that might also be related to the metaphor (i.e., sharp teeth might related to tenacity). We shall refer to the experimental property statements that were relevant to the superordinate category of the vehicle as superordinate target sentences and the experimental property statements that were relevant to only the basic-level of the vehicle as basic-level target sentences. The factorial combination of the two types of prime sentences (metaphor and literal) and the two types of target sentences (superordinate and basic-level) produced four experimental conditions, as illustrated in Table 1.

The materials also included 144 filler pairs of sentences designed to balance the number of prime and target sentences that did versus did not make sense. The filler sentence pairs were constructed similarly to the experimental pairs (i.e., the first sentence was either a metaphor or a literal sentence, and the second sentence was a property statement about the vehicle of the preceding metaphor or literal sentence). For 48 of the filler sentence pairs neither the first (metaphorical or literal) sentence nor the second (property statement) sentence made sense, for example, *Her basketball is a toll booth. Toll booths distribute hamburgers*. For 48 of the filler sentence pairs the first sentence did not make sense but the second sentence did, for example, *That apple was a tennis player. Tennis players need racquets*; and for 48 filler sentence pairs the first sentence made sense — 24 were like the experimental metaphor primes and 24 were like the experimental literal primes — and the second sentence did not make sense, for example, *My aunt's health problems are a time bomb. Bombs are special anniversary gifts, or His jacket was corduroy. Corduroy swims in the lake*. The complete set of materials for Experiment 1 is available on-line at <http://psych.wisc.edu/lang/materials/metlit.html>.

Procedure—Participants were tested in groups of four or fewer, with each participant occupying his or her own cubicle. Each cubicle contained a computer monitor and a two-key response pad. At the beginning of the session, participants read instructions on their computer monitors that explained the task and provided examples sentences. Participants were told that their task was to read each of a series of sentences and to decide rapidly and accurately whether each sentence made sense. The participants were told that some of the sentences would be metaphorical and were given examples of sentences that did and did not make sense (e.g., *My mother says that my little brother is a pig* versus *My mother says that my little brother is a desk*).

Participants were instructed to use their dominant hand to indicate their response by pressing either a key labeled “yes” or a key labeled “no.” They used their index finger to press one key and their middle finger to press the other. During the instructions, participants

responded to eight practice sentences and received feedback after each practice sentence indicating whether they responded correctly, incorrectly, or too slowly (taking longer than 5500 ms). During the actual experiment the participants did not receive this feedback. After responding to each third of the sentences in the experiment, the participants received a short break.

All sentences were presented left-justified and vertically centered so that the beginning of each sentence appeared in the same location. Each sentence remained on the screen for 5500 ms or until the participant pressed a key. Between sentences a blank screen appeared for 450 ms.

Participants—The participants were 133 community members, recruited from flyers posted around the University of Wisconsin-Madison campus and participating for monetary compensation. Only verification times for which participants responded correctly to both the target (property statement) and its preceding prime sentence were analyzed, and data from participants who did not perform better than 66% correct on each of the experimental item types were not analyzed. This criterion eliminated the data from 21 participants; the data from 112 participants were included in the analysis.

Results and Discussion

Participants' average verification times to property statements are illustrated in Figure 1. To test the hypothesis that the superordinate meaning of a vehicle is enhanced during metaphor understanding, we examined participants' average verification times to property statements that were relevant to the superordinate category. More specifically, we compared the participants' average verification times to property statements that were relevant to the superordinate category (e.g., *Sharks are tenacious*) after the participants read a metaphorical prime sentence (e.g., *That defense lawyer is a shark*), as opposed to a baseline, literal-prime sentence (e.g., *That large hammerhead is a shark*). If the metaphorical vehicle stands for the superordinate category, then we should have observed an enhancement effect. As illustrated in the two left-most bars of Figure 1, participants more rapidly verified the superordinate property statements after they read the metaphor-prime sentences ($M = 1003$ ms; $SE = 16$ ms) than after they read the literal-prime sentences ($M = 1046$ ms; $SE = 19$ ms), $F(1,111) = 7.336, p < .01$; $F(1,47) = 4.987, p < .03$.

To test the hypothesis that the basic-level meaning of a vehicle is suppressed during metaphor understanding, we examined participants' average verification times to property statements that were relevant to only the basic-level category. More specifically, we compared the participants' average verification times to property statements that were relevant to only the basic-level category (e.g., *Sharks are good swimmers*) after the participants read a metaphorical prime sentence (e.g., *That defense lawyer is a shark*), as opposed to a baseline, literal-prime sentence (e.g., *That large hammerhead is a shark*). If the metaphorical vehicle stands for the superordinate category, then we should have observed a suppression effect. As illustrated in the two right-most bars of Figure 1, participants more slowly verified the basic-level property statements after they read the metaphor-prime

sentences ($M = 1219$ ms; $SE = 22$ ms) than after they read the literal-prime sentences ($M = 1162$ ms; $SE = 18$ ms), $F(1,111) = 13.38, p < .0004$; $F(1,47) = 7.656, p < .01$.

This contrasting pattern of enhancement of the superordinate category (i.e., faster verification times for the superordinate property statements after participants read the metaphor-versus literal-prime sentences) and suppression of the basic-level meaning of the vehicle (i.e., slower verification times for the basic-level property statements after participants read the metaphor-versus literal-prime sentences) produced a statistically significant interaction, $F(1,111) = 20.27, p < .0001$; $F(1,94) = 12.29, p < .0007$.

Participants were faster to verify superordinate property statements ($M = 1024$ ms; $SE = 12$ ms) than basic-level property statements ($M = 1191$ ms; $SE = 14$ ms), $F(1,111) = 241.30, p < .0001$; $F(1,94) = 10.35, p < .002$. However, this difference was not important to our investigation and so we did not investigate its cause. Participants were very accurate to the property statements, responding incorrectly less than 2% of the time to the superordinate property statements and less than 3% of the time to the basic-level property statements. There were no differences or interactions in error rates.

The enhancement and suppression effects demonstrated by this experiment support the class inclusion theory of metaphor comprehension. Comprehending a metaphor makes properties that are central to the metaphor ground (and to the superordinate category of the vehicle) more accessible. More crucial to the theory, comprehending a metaphor makes properties of the basic-level meaning of the vehicle that are not relevant to its superordinate meaning less accessible; we suggest that they are suppressed. Such suppression of the basic-level meaning is strong evidence that metaphorical vehicles do not have basic-level meaning; they stand for the superordinate.

Although the pattern of the data from our first experiment was clear, one could argue against our interpretation of the suppression effect in the following way. Perhaps we observed that participants were slower to verify basic-level properties (e.g., *Sharks are good swimmers*) after they read metaphors (e.g., *That defense lawyer is a shark*) than after they read baseline, literal sentences (e.g., *That large hammerhead is a shark*) not because comprehending metaphors involves suppressing basic-level meanings, but because our comparison condition — our baseline, literal-prime sentences — contained the basic-level category member term (e.g., *hammerhead*). In other words, perhaps instead of demonstrating that participants were slower to verify basic-level properties after they read metaphors, we merely demonstrated that participants were faster to verify basic-level properties after they read the literal sentences, and the reason why they were faster to verify basic-level properties after they read the literal sentences is because those literal sentences contained the basic-level category word.

To rule out this explanation we conducted a second experiment that used a different baseline prime. In our second experiment our baseline prime sentences did not contain the basic-level category word (e.g., *hammerhead*); rather, the topic of the baseline prime was nonsensical when yoked with the metaphorical vehicle. For example, the baseline prime for the metaphor *His defense lawyer was a shark* was the nonsense-prime sentence, *His English*

notebook was a shark. If we are correct in our interpretation of Experiment 1, then the nonsense baseline prime in Experiment 2 should yield exactly the same results as we observed in Experiment 1.

EXPERIMENT 2

Methods

Materials—Experiment 2 used the same 48 metaphor-prime sentences as Experiment 1, with occasional minor modifications (e.g., change of a demonstrative reference to a pronoun). For each metaphor-prime sentence we constructed a nonsense-prime sentence by changing the topic of the corresponding (sensible) metaphor-prime sentence. For example, for the metaphor-prime sentence, *His defense lawyer is a shark*, we replaced *His defense lawyer* with *His English notebook* making the nonsense-prime sentence, *His English notebook is a shark*.

As in Experiment 1, each prime sentence was followed by a target sentence, which was a property statement. The target sentences in Experiment 2 were identical to those of Experiment 1. For the experimental target sentences, the property was relevant to the superordinate category represented by the vehicle of the preceding metaphor- or nonsense-prime sentence, or the property was relevant to only the basic-level category represented by the vehicle of the preceding metaphor-prime or nonsense-prime sentence. The factorial combination of the two types of prime sentences (metaphor versus nonsense) and the two types of target sentences (property statements relevant to the superordinate versus basic-level) produced four experimental conditions, as illustrated in Table 2.

The materials also included 48 filler pairs of sentences designed to balance the number of prime and target sentences that did versus did not make sense. The filler sentence pairs were constructed similarly to the experimental pairs, in that the first sentence was either a metaphor or a nonsense statement, and the second sentence was a property statement about the vehicle of the preceding metaphor or nonsense statement. For 24 of the filler sentence pairs neither sentence made sense, for example, *Her basketball is a toll booth*. Toll booths distribute hamburgers. For the other 24 filler sentence pairs, the first sentence made sense, but the second sentence did not make sense, for example, *My aunt's health problems are a time bomb*. Time bombs are special wedding gifts. The complete set of materials for Experiment 2 is available on-line at <http://psych.wisc.edu/lang/materials/metnon.html>.

Procedure—The procedure was identical to that of Experiment 1.

Participants—One hundred and ninety one undergraduates at the University of Wisconsin-Madison participated for extra credit in an introductory psychology course. All participants were native American English speakers. As in Experiment 1, only verification times for which participants responded correctly to both the target (property statement) and its preceding prime sentence were analyzed, and data from participants who failed to perform better than 66% correct on each of the experimental item types were not analyzed. This criterion eliminated the data from 51 participants; the data from 140 participants were included in the analysis.

Results and Discussion

Participants' average verification times to property statements are illustrated in Figure 2. To test the hypothesis that the superordinate meaning of a vehicle is enhanced during metaphor understanding, we examined participants' average verification times to property statements that were relevant to the superordinate category. More specifically, we compared the participants' average verification times to property statements that were relevant to the superordinate category (e.g., *Sharks are tenacious*) after the participants read a metaphor-prime sentence (e.g., *His defense lawyer is a shark*), as opposed to a baseline, nonsense-prime sentence (e.g., *His English notebook is a shark*). If the metaphorical vehicle stands for the superordinate category, then we should have observed an enhancement effect. And again, as illustrated in the two left-most bars of Figure 2, participants more rapidly verified the superordinate property statements after they read the metaphor-prime sentences ($M = 1093$ ms; $SE = 19$ ms) than after they read the nonsense-prime sentences ($M = 1173$ ms; $SE = 20$ ms), $F(1,139) = 20.22, p < .0001$; $F(2,1,47) = 21.20, p < .0001$.

To test the hypothesis that the basic-level meaning of a vehicle is suppressed during metaphor understanding, we examined participants' average verification times to property statements that were relevant to only the basic-level category. More specifically, we compared the participants' average verification times to property statements that were relevant to only the basic-level category (e.g., *Sharks are good swimmers*) after the participants read a metaphor-prime sentence (e.g., *His defense lawyer is a shark*), as opposed to a baseline, nonsense-prime sentence (e.g., *His English notebook is a shark*). If the metaphorical vehicle stands for the superordinate category, then we should have observed a suppression effect. And again, as illustrated in the two right-most bars of Figure 2, participants more slowly verified the basic-level property statements after they read the metaphor-prime sentences ($M = 1341$ ms; $SE = 22$ ms) than after they read the nonsense-prime sentences ($M = 1273$ ms; $SE = 20$ ms), $F(1,139) = 14.46, p < .0002$; $F(2,1,47) = 4.26, p < .05$.

As in Experiment 1, the contrasting pattern of enhancement of the superordinate category (i.e., faster verification times for the superordinate property statements after participants read the metaphor-prime versus nonsense-primers) and suppression of the basic-level meaning of the vehicle (i.e., slower verification times for the basic-level property statements after participants read the metaphor-prime versus nonsense-primers) produced a statistically significant interaction, $F(1,139) = 34.43, p < .0001$; $F(2,1,94) = 18.15, p < .0001$.

Participants were again faster to verify superordinate property statements ($M = 1133$ ms; $SE = 14$ ms) than basic-level property statements ($M = 1307$ ms; $SE = 15$ ms), $F(1,139) = 174.69, p < .0001$; $F(2,1,94) = 9.44, p < .003$. However, this difference was not important to our investigation and so we did not investigate its cause. Participants responded incorrectly less than 1% of the time to the superordinate property statements and less than 2% of the time to the basic-level property statements. There were no differences or interactions in error rates.

Experiment 2 both replicated the effects found in Experiment 1 and allowed us to reject the alternative explanation to the suppression effect. Recall that the alternative explanation

assumed that the suppression effect found in Experiment 1 was actually an enhancement following the literal prime sentence. It suggested that the presence of a basic-level concept (e.g., *hammerhead*) in the literal-prime sentence might have facilitated responses to the basic-level property, making it seem like the metaphor-prime suppressed responses to that property. The fact that Experiment 2 yielded the same result with a baseline prime that did not use the basic-level concept argues against this alternative.

EXPERIMENT 3

We conducted the second experiment in order to rule out an alternative explanation to Experiment 1's suppression effect. Though Experiment 2, even with a nonsensical baseline, showed the same suppression pattern as Experiment 1, a critic might still argue that the nonsense prime is not a "truly" neutral baseline. It is possible to explain the suppression effect of Experiment 2 by assuming that the nonsense prime, *His English notebook is a shark* primed *good swimmers*, speeding response time to the target sentence, *Sharks are good swimmers*. Perhaps the word *shark* in the nonsensical string is understood at the basic-level of categorization, thus enhancing the basic-level properties of *shark*, and consequently facilitating the comprehension of the target sentence *Sharks are good swimmers*. In addition, one could assume that the metaphor, *My lawyer is a shark* does not affect the availability of the property "are good swimmers" at all. If this is the case, then readers should be slower to respond to the target following the metaphor than the nonsense prime. Thus, what looks like a suppression effect in Experiment 2 could simply be an enhancement effect.

We designed Experiment 3 to overcome this criticism. Ideally, one would want to use a truly neutral priming sentence. Unfortunately, such neutral primes do not exist because any linguistic material has the potential to prime some aspect of the target sentence. We therefore aimed at addressing the specific problem that the nonsense prime poses. The heart of the alternative explanation is that the word *shark* in the nonsense prime might have enhanced the comprehension of the target sentence. Therefore, the baseline prime in Experiment 3 did not include the vehicle name from the corresponding metaphor prime. Instead of a nonsense prime, this experiment used an unrelated metaphor as a baseline prime. For example, the unrelated (baseline) prime for the target sentence, *Sharks are good swimmers* was *That new student is a clown*. An unrelated metaphor cannot be expected to systematically induce enhancement of the basic-level property such as *are good swimmers*. Therefore, with this unrelated metaphor baseline, a suppression effect can no longer be explained in terms of enhancement.

But the new baseline solved one problem and created another. The metaphor prime and the metaphor baseline are no longer equated in the most important way. Only the metaphor prime includes the crucial term, the vehicle *shark*. In this new configuration, the word *shark* is repeated between the metaphor primes and the target sentence, but not between the unrelated baseline primes and the target sentences. It is well established that readers are faster to identify words they have seen shortly before, so it is virtually guaranteed that the repetition of *shark* alone would result in faster response time following the metaphor than the baseline prime. This should be true even for the responses to the basic-level target sentences (*Sharks are good swimmers*).

Though such a repetition priming effect is bound to occur, it is irrelevant to our theoretical concern and would add only a constant in the metaphor prime conditions. To avoid this potential problem, we will standardize the verification latencies within each prime type. Any remaining differences should more directly reflect enhancement and suppression effects.

Methods

Materials—We paired each metaphor prime sentence (e.g. *That defense lawyer is a shark*) with another unrelated metaphor prime sentence (e.g. *That new student is a clown*) so that the same metaphor served as both a relevant prime in one sentence set and an unrelated prime in another. As Table 3 illustrates, this simple pairing created completely counterbalanced sentence sets, with unrelated metaphors serving as the baseline primes. The complete set of materials is available online at <http://psych.wisc.edu/lang/materials/metneut.html>

Procedure—The procedure was identical to that of Experiments 1 and 2.

Participants—One hundred and sixty-seven undergraduates at the University of Wisconsin-Madison participated for extra credit in an introductory psychology course. All participants were native American English speakers. As in Experiments 1 and 2, only verification times for which participants responded correctly to both the target (property statement) and its preceding prime sentence were analyzed. Data from participants who failed to perform better than 66% correct on each of the experimental item types were not analyzed. This criterion eliminated the data from 51 of the total participants; the data from 116 participants were included in the analysis.

Results and Discussion

Participants' average verification times to property statements are illustrated in Figure 3. As in Experiments 1 and 2, we observed a statistically significant interaction, $F(1,155) = 6.648, p < .01$; $F(2,147) = 4.046, p < .05$. However, as seen in Figure 3, this interaction is not of the 'cross-over' style observed in Experiments 1 and 2, most likely because the lack of word repetition inflated the verification times to target sentences following the unrelated primes. Indeed, if one were to subtract an estimated penalty for not having a repeated word, say 260 ms, from the average verification times to property statements following the unrelated primes (those that did not contain a repeated word), the signature cross-over interaction would be perfectly manifested. However, a more elegant solution is to standardize the participants' verification latencies within each prime type.

Figure 4 presents the standardized verification latencies for Experiments 1, 2, and 3. As Figure 4 illustrates, when we standardized the verification latencies within prime types, the pattern of suppression and enhancement that we observed in Experiments 1 and 2 appears in Experiment 3. In all three experiments, participants were faster to verify the superordinate property statements after reading the metaphor-prime sentences (E1 $M = -.473, SE = .07$; E2 $M = -.460, SE = .07$; E3 $M = -.454, SE = .07$) than after reading the unrelated-prime sentences (E1 $M = -.287, SE = .09$; E2 $M = -.206, SE = .08$; E3 $M = -.276, SE = .07$), E1

$F(1, 112) = 6.67, p < .01$; $E2 F(1, 139) = 13.78, p < .0001$; $E3 F(1,155) = 7.31, p < .01$. This pattern demonstrates support for our enhancement hypothesis.

Conversely, in all three experiments, participants were also slower to verify the basic-level property statements after reading the metaphor-prime sentences ($E1 M = .473, SE = .09$; $E2 M = .460, SE = .08$; $E3 M = .454, SE = .07$) than after reading the unrelated-prime sentences ($E1 M = .287, SE = .09$; $E2 M = .206, SE = .08$; $E3 M = .276, SE = .08$), $E1 F(1, 112) = 6.67, p < .01$; $E2 F(1, 139) = 13.78, p < .0001$; $E3 F(1,155) = 7.31, p < .01$. For each experiment, this contrasting pattern led to a reliable interaction, $E1 F(1, 112) = 13.34, p < .001$; $E2 F(1, 139) = 27.55, p < .0001$; $E3 F(1,155) = 14.621, p < .001$, which when computed on the standardized verification latencies was always of the ‘cross-over’ type, as shown in Figure 4.

GENERAL DISCUSSION

The experiments reported here provide two central findings. They consistently demonstrate an enhancement effect: Properties of a metaphor’s vehicle that are central to the metaphorical meaning become more accessible. These experiments also demonstrate an interesting suppression effect: Properties of a metaphor’s vehicle that are relevant only to the basic-level meaning of the vehicle, and are not relevant to the superordinate meaning of the vehicle, become less accessible. We suggest that these properties are suppressed — to enable metaphor understanding. In the same way, inappropriate meanings of polysemous words are suppressed to enable lexical access, previously mentioned concepts that are not the antecedents of an anaphor are suppressed to enable anaphoric reference, competing topics are suppressed to enable cataphoric reference, and so forth. The enhancement and suppression effects occurred with three different baselines, a “literal” baseline, a “nonsense” baseline, and a “neutral” (unrelated) baseline, suggesting that the phenomena are robust.

Implications for Theories of Metaphor

The class inclusion theory of metaphor (Glucksberg & Keysar, 1990) predicts precisely the effects of suppression and enhancement that we observed in our two experiments. The class inclusion theory assumes that metaphors such as *cigarettes are time-bombs* are understood as is — as category statements. *Time-bomb* stands for a superordinate category, not for the object *time-bomb*. The enhancement effect suggests that indeed the superordinate category is accessible during metaphor understanding. The suppression effect more directly demonstrates that the vehicle does not stand for the basic-level meaning. The fact that basic-level properties were less accessible strongly suggests that the vehicle does not stand for the object-level category. These experiments, together with experiments by Glucksberg, McGlone and Manfredi (1997), converge on direct empirical evidence for the class inclusion theory of metaphor.

While the experiments provide direct support for the class inclusion theory because they verify its natural predictions, we do not claim that they refute competing theories of metaphor understanding. As with all experiments to date in the study of metaphor understanding, these are not “critical experiments.” The reason is, that our experiments do not directly test the predictions of alternative theories. For example, a central assumption in

Wolff and Gentner's (1992) theory is that metaphor understanding involves an early comparison between the topic and the vehicle. Our data do not speak to this issue either way. We cannot tell from the procedure and the results whether such a comparison takes place. While the experiments do not directly refute alternative theories they provide strong supporting evidence to the class inclusion theory, and they need to be explained by future theoretical attempts.

Category Mistake, Category Retrieval, or Category Construction?

One of the persisting puzzles in the empirical study of metaphor has been that while a metaphorical sentence is, strictly speaking, anomalous (because it involves a category mistake, Davidson, 1978), experiments consistently find no evidence for the psychological reality of such an anomaly. It seems like readers and addressees comprehend a metaphor with no sense of an initial aberration. The class inclusion theory and our results provide one resolution to this puzzle. If our interpretation of the results is correct, then readers need not encounter an anomaly when they read the sentence *His defense lawyer is a shark*. The reason is, that once *shark* stands for the superordinate ad-hoc category, the sentence is understood as a normal category statement. To arrive at the appropriate meaning of *shark* comprehenders must do some "work," just like they work at identifying and constructing the context-dependent meaning of words in general.

The class inclusion account, then, explains the suppression effect as a result of the construction of the ad-hoc category that is the basis for understanding the metaphor. One might suggest instead that our results reflect no category construction but category retrieval. Take for example an extreme case, a highly conventionalized metaphorical use of *butcher* as in *his surgeon is a butcher*. The metaphorical meaning of *butcher* is so lexicalized in English that it appears as an entry in several dictionaries. For example, it appears in the American Heritage Dictionary as the fourth entry "one who bungles something." If by analogy such a meaning of *butcher* is also an entry in our mental lexicon, then it is possible that lexical ambiguity resolution is taking place. In that case, the suppression effect might reflect the same process that underlies lexical access of polysemous words such as *spade* as Gernsbacher and her colleagues have demonstrated (e.g., Gernsbacher & Faust, 1991b). If so, the suppression effect for *butcher* should be taken to reflect not the construction of an ad-hoc category but the retrieval of the pre-existing lexicalized category of people who bungle things.

We propose that this alternative interpretation of our data should be rejected because it makes two false predictions. The first prediction concerns a possible priming effect. If the metaphorical meaning of a vehicle is really lexicalized then the vehicle by itself should be an effective prime for terms that are related to the metaphorical meaning. However, experiments by Blasko and Connine (1993) suggest the opposite. While they found that conventional metaphors prime the metaphorical meaning very early in the process, they found no such priming by the conventional vehicles in isolation. Therefore, a "lexical access-type" interpretation of our results could be rejected on these grounds.

A stronger reason to reject this alternative interpretation comes from further analysis of our own data. We reasoned that if the suppression effect reflects lexical access for highly

conventionalized metaphors, then the effect should correlate positively with ease of understanding the metaphor. We tested this possibility with the data from Experiment 2. We operationalized metaphor conventionality as error rate to the metaphor primes, assuming that the more conventional the metaphor the fewer errors participants should make in response to it. (Recall that the participants' task was to decide whether each sentence made sense; therefore, saying that one of our metaphor-prime sentences did not make sense would be considered an error.) We then computed a suppression score (an effect size) for each metaphor by subtracting (a) the mean verification time to its basic-level relevant property statement when that property statement was preceded by the nonsense-prime from (b) the mean verification time to its basic-level relevant property statement when that property statement was preceded by the metaphor-prime. Although error rates ranged across metaphor primes from 0% (*Her husband is a gem*) to 87% (*The baby monkey is a vine*), there was no hint of correlation between these error rates and the suppression scores ($r = .002$). We therefore feel confident rejecting this alternative explanation and proposing instead that our results reflect the process of category construction.

From Stage Models to Levels of Categorization

One might want to interpret our findings radically differently, in a way that supports the traditional stage model of metaphor understanding as proposed by Searle (1993). Recall that the traditional model proposes three stages of metaphor understanding: Initially people interpret the sentence literally; then they compare it with context; if they encounter a discrepancy, they interpret it metaphorically. One might want to use this model to explain the suppression effect. Perhaps just because participants construct a literal interpretation of the metaphor, once they reject it in favor of the metaphorical interpretation they must suppress it. According to this model, our data would support a literal-first stage in metaphor understanding.

However, such an interpretation of our data would be misguided; it confuses sentence meaning with word meaning. To support Searle's stage theory one must show that metaphor comprehension is preceded by the construction of a literal, compositional meaning of the metaphorical sentence. It is not sufficient to show that the lexical meaning of the sentence's words was accessed. For example, one must show that the first stage of understanding *his defense lawyer is a shark* involves the construction of a proposition where *his defense lawyer* is categorized as a type of fish. While our data show evidence for lexical access of the vehicle, it provides no evidence that the literal, compositional meaning of the metaphorical sentence was actually computed.

In contrast, we interpret our data as suggesting that the basic-level meaning of the vehicle is activated. The proposal that basic-level meaning is available during comprehension is not surprising; its eventual suppression is the interesting fact. The reason that the availability of the basic-level meaning of the vehicle is not surprising is that we know that in general accessing word meaning is not optional (Stroop, 1935), and that even contextually inappropriate meanings of words might be activated initially, even in a strongly biasing context (Gernsbacher & Faust, 1991b; Simpson, 1981; Swinney, 1979). Research on figurative language demonstrates the same availability of literal word meaning. For

example, Blasko and Connine (1993) demonstrated that literal associates of a metaphorical vehicle are activated following metaphor understanding, and Cacciari and Tabossi (1988) showed that even with non-compositional phrases such as idioms, there is activation of the literal associates of the final word of the idiom. If the vehicle's basic-level meaning is indeed available during comprehension, does it play any role and why is it suppressed?

There is one interesting difference between the availability of the basic-level meaning of the vehicle and the availability of a contextually inappropriate meaning of a polysemous word. When, for example, the *money*-related meaning of *bank* is activated — even when the context is clearly about the bank of a river (Swinney, 1979) — it is reasonable to assume that this inappropriate meaning does not play any role in understanding the intended meaning. In contrast, we assume that the availability of the basic-level meaning of the vehicle is functional. The comprehender might be constructing an ad-hoc category using the basic-level meaning as a prototype in order to arrive at the superordinate meaning. In this sense, the basic-level meaning plays a role in the very understanding of the metaphor. But once the superordinate meaning is constructed, the basic-level meaning is potentially confusing. It might be confusing because if the metaphor is understood as class inclusion, and the basic-level meaning of the vehicle is used, then a category mistake or an anomaly should result. Therefore, to allow a straightforward understanding of the metaphor as an unqualified category statement, the basic-level meaning of the vehicle is suppressed -- for the same reason that the irrelevant meaning of a polysemous word is suppressed.

The suppression effect also bears on a related issue raised in Glucksberg and Keysar (1990). The class inclusion theory of metaphor suggested that the vehicle has a dual reference function: to refer at both the basic-level and the superordinate-level of categorization. For example, the theory assumed that the vehicle *shark* refers simultaneously to the fish and to the ad-hoc category of tenacious, vicious things that it typifies. The intuition behind that idea was that such dual reference would allow us to use both the categorization form *his defense lawyer is a shark* and the comparison form *his defense lawyer is like a shark*. Kittay (1997) points out that the idea of dual reference in the study of metaphor appears in different forms in the writings of a variety of metaphor researchers (e.g., Henle, 1965; Kittay, 1987; Pierce, 1931). In contrast to this tradition, the suppression effect that we report here suggests that perhaps metaphor understanding does not involve dual reference. Instead, the vehicle uniquely refers to the superordinate category. While language users can reflect about a metaphor and realize that the vehicle can have a dual reference, it is possible that to understand the metaphor they suppress one of these referents so that they can more naturally comprehend it as a true category statement.

Comprehension Mechanisms and Metaphor Understanding

Not only does the understanding of metaphor not involve anomaly, but experimental research shows that metaphor understanding is as natural as the understanding of literal language. Our paper suggests one central reason for the comparability of metaphor and literal comprehension: They are understood via exactly the same general comprehension mechanisms. Specifically, we demonstrated that the same mechanisms of enhancement and suppression that underlie comprehension in general (e.g., Gernsbacher & Faust, 1991a) and

language comprehension in particular (e.g., Gernsbacher & Jescheniak, 1995;), are also crucial for metaphor understanding. Our findings suggest that metaphor comprehension could be naturally accounted for within the Structure Building Framework (Gernsbacher, 1990; 1991a; 1995; 1997c).

These general comprehension mechanisms support both metaphor and literal comprehension in two important ways. First, the understanding of a metaphorical meaning, as we demonstrated in our experiments, is comparable to any other like-structure literal sentence. Enhancement and suppression enable understanding of a nominative metaphor as they support the understanding of any other categorization statement. Second, in earlier work Keysar (1994) showed that suppression is involved in inferences regarding intended meaning, and that it operates the same for literal and metaphorical meaning. Those studies demonstrated that suppression plays a role in arriving at a final intended meaning of sentences in discourse context: Readers can infer a metaphorical intended meaning in context by suppressing a literal meaning, but they can also infer a literal intended meaning by suppressing a metaphorical meaning.

CONCLUSION

On the one hand, our experiments provide support for the class inclusion theory of metaphor understanding. On the other hand, the experiments demonstrate the way that general comprehension mechanisms, as described by Gernsbacher and her colleagues, work at the service of metaphor understanding. Most importantly, the suppression mechanism is crucial in eliminating potentially confusing information during metaphor interpretation, just like it is important in suppressing irrelevant information during any act of understanding.

Acknowledgments

This research was supported by grants from the National Institutes of Health (RO1 NS 29926) and the Army Research Institute (DASW0194-K-0004, DASW0196-K-0013, and DAAG55-97-1-0224) awarded to Morton Ann Gernsbacher. We are grateful to Jennifer Deaton, Kathryn McEwen, and Nitika Gondal for assisting with the early stages of this research, and to Lisa Bach, Vaughn Brandt, Karla Baumann, Mary Campana, Karin Christoph, Shawna Clark, Fred Dick, Laurie Loomans, Martha Kralovec, Alexis Halaby, Brenda Hallada, Erica Herdey, Amelia Nearing, Jennifer Nelson, Lana Nenide and Melissa Rosenkranz for testing the participants.

REFERENCES

- Altwerger, B.; Strauss, SL. Readers' understanding of metaphor: Implications of linguistic theory for psycholinguistic research. In: Tierney, R.; Anders, P.; Mitchell, J., editors. *Understanding readers' understanding*. Erlbaum; Hillsdale, NJ: 1987. p. 175-188.
- Barsalou L. Ad-hoc categories. *Memory & Cognition*. 1983; 11:211–227. [PubMed: 6621337]
- Bassok, M. Two types of reliance on correlations between content and structure in reasoning about word problems. In: English, L., editor. *Mathematical reasoning: Analogies, metaphors, and images*. Erlbaum; Hillsdale, NJ: in press
- Blasko DG, Connine CM. Effects of familiarity and aptness on metaphor processing. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 1993; 2:295–308.
- Brown, R. *Words and things*. The Free Press; New York: 1958.
- Cacciari C, Tabossi P. The comprehension of idioms. *Journal of Memory and Language*. 1988; 27:668–683.

- Carbonell, JG. Derivational analogy: A theory of reconstructive problem solving and expertise acquisition. In: Michalski, RS.; Carbonell, JG.; Mitchell, TM., editors. *Machine learning*. Tioga Publishing Co.; Palo Alto: CA: 1986. p. 371-392.
- Clement CA, Gentner D. Systematicity as a selection constraint in analogical mapping. *Cognitive Science*. 1991; 15:89–132.
- Davidson D. What metaphors mean. *Critical Inquiry*. 1978; 5:31–47.
- Gentner, D.; Falkenhainer, B.; Skorstand, J. Viewing metaphor as analogy. In: Helman, DH., editor. *Analogical reasoning: perspectives of artificial intelligence, cognitive science, and philosophy*. Kluwer Academic Publishers; Boston, MA: 1988. p. 171-177.
- Gernsbacher MA. Mechanisms that improve referential access. *Cognition*. 1989; 32:99–156. [PubMed: 2752708]
- Gernsbacher, MA. *Language comprehension as structure building*. Erlbaum; Hillsdale, NJ: 1990.
- Gernsbacher, MA. Cognitive processes and mechanisms in language comprehension: The structure building framework. In: Bower, GH., editor. *The psychology of learning and motivation*. Academic Press; New York: 1991a. p. 217-263.
- Gernsbacher MA. Comprehending conceptual anaphors. *Language and Cognitive Processes*. 1991b; 6:81–105.
- Gernsbacher MA. Less skilled readers have less efficient suppression mechanisms. *Psychological Science*. 1993; 4:294–298.
- Gernsbacher, MA. The costs and benefits of meaning; Paper presented at the 35th annual meeting of the Psychonomics Society; St. Louis, MO. 1994; Nov.
- Gernsbacher, MA. The Structure Building Framework: What it is, what it might also be, and why. In: Britton, BK.; Graesser, AC., editors. *Models of text understanding*. Erlbaum; Hillsdale, NJ: 1995. p. 289-311.
- Gernsbacher MA. Generic pronominal anaphora: The case of the English singular *they*. *Verbum*. 1997a; 19:67–84.
- Gernsbacher MA. Group differences in suppression skill. *Aging, Neuropsychology, and Cognition*. 1997b; 4:175–184.
- Gernsbacher MA. Two decades of structure building. *Discourse Processes*. 1997c; 23:265–304.
- Gernsbacher MA, Faust ME. The mechanism of suppression: A component of general comprehension skill. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 1991a; 17:245–262.
- Gernsbacher, MA.; Faust, ME. The role of suppression in sentence comprehension. In: Simpson, GB., editor. *Understanding word and sentence*. Elsevier Science Publishers; 1991b. p. 97-128.
- Gernsbacher, MA.; Faust, M. Skilled suppression. In: Dempster, FN.; Brainerd, CN., editors. *New perspectives on interference and inhibition in cognition*. Academic Press; San Diego, CA: 1994. p. 295-327.
- Gernsbacher MA, Jescheniak JD. Cataphoric devices in spoken discourse. *Cognitive Psychology*. 1995; 29:24–58. [PubMed: 7641525]
- Gernsbacher MA, Robertson RRW. Reading skill and suppression revisited. *Psychological Science*. 1995; 6:165–169.
- Gernsbacher MA, Shroyer S. The cataphoric use of the indefinite *this* in spoken narratives. *Memory & Cognition*. 1989; 17:536–540. [PubMed: 2796738]
- Gernsbacher, MA.; St. John, MF. Modeling the mechanism of suppression in lexical access. In: Gorfein, DS., editor. *On the consequences of meaning*. in press
- Gernsbacher MA, Varner KR, Faust ME. Investigating differences in general comprehension skill. *Journal of Experimental Psychology: Learning, Memory and Cognition*. 1990; 16:430–445.
- Gerrig RJ. Empirical constraints on computational theories of metaphor: Comments on Indurkha. *Cognitive Science*. 1989; 13:235–241.
- Gerrig RJ, Healy AF. Dual processes in metaphor understanding: Comprehension and appreciation. *Journal of Experimental Psychology: Learning, Memory and Cognition*. 1983; 9:667–675.
- Gibbs RW. Literal meaning and psychological theory. *Cognitive Science*. 1984; 8:275–304.
- Gick ML, Holyoak KJ. Schema induction and analogical transfer. *Cognitive Psychology*. 1983; 15:1–38.

- Glucksberg S. Beyond literal meaning: The psychology of allusion. *Psychological Science*. 1991; 2:146–152.
- Glucksberg S. Understanding metaphors. *Current Directions in Psychological Science*. 1998; 7:39–43.
- Glucksberg S, Gildea P, Bookin H. On understanding nonliteral speech: Can people ignore metaphors? *Journal of Verbal Learning and Verbal Behavior*. 1982; 21:85–98.
- Glucksberg S, Keysar B. Understanding metaphorical comparisons: Beyond similarity. *Psychological Review*. 1990; 97:3–18.
- Glucksberg, S.; Keysar, B. How metaphors work. In: Ortony, A., editor. *Metaphor and thought*. Second Edition. Cambridge University Press; Cambridge: 1993. p. 401–424.
- Glucksberg, S.; Manfredi, DA.; McGlone, MS. Metaphor comprehension: How metaphors create new categories. In: Ward, TB.; Smith, SM.; Vaid, J., editors. *Creative thought: An investigation of conceptual structures and processes*. American Psychological Association; Washington, DC: 1997. p. 327–350.
- Glucksberg S, McGlone MS. When love is not a journey: What metaphors mean. *Journal of Pragmatics*. 1999; 31:1541–1558.
- Glucksberg S, McGlone MS, Manfredi D. Property attribution in metaphor comprehension. *Journal of Memory and Language*. 1997; 36:50–67.
- Grice, HP. Logic and conversation. In: Cole, P.; Morgan, J., editors. *Syntax and semantics 3: Speech acts*. Academic Press; New York: 1975. p. 41–58.
- Hage P, Miller WR. ‘Eagle’ = ‘bird’: A note on the structure and evolution of Shoshoni ethnoornithological nomenclature. *American Ethnologist*. 1976; 3:481–488.
- Henle, P. Metaphor. In: Henle, P., editor. *Language, thought and culture*. Ann Arbor Paperbacks; Ann Arbor, MI: 1965. p. 173–195.
- Holyoak KJ, Thagard P. Analogical mapping by constraint satisfaction. *Cognitive Science*. 1989; 13:295–355.
- Keysar B. On the functional equivalence of literal and metaphorical interpretations in discourse. *Journal of Memory and Language*. 1989; 28:375–385.
- Keysar B. Discourse context effects: Metaphorical and literal interpretations. *Discourse Processes*. 1994; 18:247–269.
- Kintsch W. Metaphor comprehension: A computational theory. *Psychonomic Bulletin & Review*. 2000; 7:257–266. [PubMed: 10909133]
- Kittay, EF. *Metaphor*. Oxford University Press; New York: 1987.
- Kittay, EF. Of “men” and metaphors: Shakespeare, embodiment, and filing cabinets. In: Ward, TB.; Smith, SM.; Vaid, J., editors. *Creative thought: An investigation of conceptual structures and processes*. APA; Washington, DC: 1997. p. 375–402.
- Levin, SR. *The semantics of metaphor*. The Johns Hopkins University Press; Baltimore: 1977.
- Malgady R, Johnson M. Modifiers in metaphor: Effects of constituent phrase similarity on the interpretation of figurative sentences. *Journal of Psycholinguistic Research*. 1976; 5:43–52.
- Marschark M, Katz A, Paivio A. Dimensions of metaphor. *Journal of Psycholinguistic Research*. 1983; 12:17–40.
- Marschark M, Reed Hunt R. On memory for metaphor. *Memory & Cognition*. 1985; 13:413–424. [PubMed: 4088051]
- McGlone MS, Manfredi DA. Topic vehicle-interaction in metaphor comprehension. *Memory & Cognition*. in press.
- Miller, GA. Images and models: Similes and metaphors. In: Ortony, A., editor. *Metaphor and thought*. Second Edition. Cambridge University Press; Cambridge: 1993. p. 357–400.
- Newport, EL.; Bellugi, U. Linguistic expressions of category levels in a visual-gesture language: A flower is a flower is a flower. In: Rosch, E.; Lloyd, BB., editors. *Cognition and categorization*. Erlbaum; Hillsdale, NJ: 1978.
- Ortony A. Beyond literal similarity. *Psychological Review*. 1979; 86:161–180.
- Ortony A, Schallert DL, Reynolds RE, Antos SJ. Interpreting metaphors and idioms: Some effects of context on comprehension. *Journal of Verbal Learning and Verbal Behavior*. 1978; 17:465–477.
- Pierce, CS. *Collected papers*. Harvard University Press; Cambridge, MA: 1931.

- Searle, J. Metaphor. Metaphor and thought. 2nd Edition. Ortony, A., editor. Cambridge University Press; Cambridge: 1993. p. 83-111.
- Simpson GB. Meaning dominance and semantic context in the processing of lexical ambiguity. *Journal of Verbal Learning and Verbal Behavior*. 1981; 20:120–136.
- Sternberg RJ. Component processes in analogical reasoning. *Psychological Review*. 1977; 84:353–378.
- Stroop JR. Studies of interference in serial verbal reactions. *Journal of Experimental Psychology*. 1935; 18:643–662.
- Swinney DA. Lexical access during sentence comprehension: (Re)consideration of context effects. *Journal of Verbal Learning and Verbal Behavior*. 1979; 18:645–659.
- Thagard, P. Dimensions of analogy. In: Helman, DH., editor. *Analogical reasoning: perspectives of artificial intelligence, cognitive science, and philosophy*. Kluwer Academic Publishers; Boston, MA: 1988. p. 105-124.
- Tourangeau R, Sternberg RJ. Aptness in metaphor. *Cognitive Psychology*. 1981; 13:27–55.
- Tourangeau R, Sternberg RJ. Understanding and appreciating metaphors. *Cognition*. 1982; 11:203–244. [PubMed: 7199412]
- Wolff, P.; Gentner, D. Proceedings of the 14th annual Cognitive Science Society. Erlbaum; Hillsdale, NJ: 1992. The time course of metaphor comprehension; p. 504-509.




Figure 1. Participants' average verification times for property statements in Experiment 1.




Figure 2.
Participants' average verification times for property statements in Experiment 2.




Figure 3.
Participants' average verification times for property statements in Experiment 3.




Figure 4.
Standardized verification times for property statements in Experiments 1, 2, and 3.

Table 1
Example Experimental Stimuli for Experiment 1

		Prime Sentence	
		Metaphor	Literal
Target	Superordinate Relevant	That defense lawyer is a shark. Sharks are tenacious.	That large hammerhead is a shark. Sharks are tenacious.
Sentence	Basic-Level Relevant	That defense lawyer is a shark. Sharks are good swimmers.	That large hammerhead is a shark. Sharks are good swimmers.

Table 2
Example Experimental Stimuli for Experiment 2

		Prime Sentence	
		Metaphor	Nonsense
Target	Superordinate Relevant	His defense lawyer is a shark. Sharks are tenacious.	His English notebook is a shark. Sharks are tenacious.
Sentence	Basic-Level Relevant	His defense lawyer is a shark. Sharks are good swimmers.	His English notebook is a shark. Sharks are good swimmers.

Table 3
Example Experimental Stimuli for Experiment 3

		Prime Sentence	
		Metaphor	Unrelated
Target	Superordinate Relevant	That defense lawyer is a shark. Sharks are tenacious.	That new student is a clown. Sharks are tenacious.
Sentence	Basic-Level Relevant	That defense lawyer is a shark. Sharks are good swimmers.	That new student is a clown. Sharks are good swimmers.