Nonliteral Language Processing

7

Unless you are educated in metaphor, you are not safe to be let loose in the world.

ROBERT FROST

Thus far, we have been treating language as though there was a straightforward relationship between what the speaker said and what the speaker meant. Now consider the following expressions:

Can you open the door?
He's a real stud.
The stop light went from green to red.

What do these expressions have in common? What they have in common is that what they say (the *standard* or *literal meaning*) and what they mean (the *speaker* or *intended meaning*) are different.

Compare *Can you open the door?* to *Can you bench press 200 pounds?* On the surface, these two expressions look very similar, but they are interpreted in very different ways. Most people interpret the first as a polite way of saying, *I want you to open the door*. Most people interpret the second differently, not as a polite way of asking someone to lift 200 pounds, but rather as a direct request for information about how strong the addressee is. If a speaker wanted to do the polite request form of the bench-pressing question, they would most likely say something like *Please bench press 200 pounds. I would like to watch*. Similarly, under all but the most limited circumstances, *He's a real stud* does not imply that the speaker thinks the subject is a male horse. Instead, the speaker is expressing the opinion that the subject possesses the qualities of a virile, strong male person.

Introduction to Psycholinguistics: Understanding Language Science, Second Edition.
Matthew J. Traxler.

Companion Website: www.wiley.com/go/traxler/psycholinguistics2e

But doesn't *The stop light went from green to red* mean exactly that "The stop light went from green to red"? The issue here involves the verb *went*. *Went* is the past tense form of the verb go, whose default, literal meaning expresses motion from one place to another. So, you might say *The plane went from Los Angeles to Sacramento* to express motion of the plane along a pathway. But stop lights don't move along a pathway. They are stationary. What actually happened is that one light stopped shining right before a different light started shining.[1] Nothing moved, but we still use a *verb of motion* to describe the event.

In all three of these cases, we are dealing with *nonliteral* language, because there is a difference between the standard meaning, based on the way the words in the expression are normally interpreted, and the meaning a listener actually assigns to the utterance. Nonliteral language creates a number of challenges for the language comprehension system, because the relationship between what is said and what is meant is not always easy or straightforward for the listener to discover. This chapter reviews different types of nonliteral language, explains why nonliteral expressions create challenges for any interpretation system, and describes different theories about how nonliteral language is processed and interpreted.

Types of Nonliteral Language

Nonliteral language comes in several different forms. There are *indirect requests*, like *Can you open the door? Do you have the time? Would you pass the salt?* (which contrast with direct requests or commands: *Open the door! Tell me the time! Give me some salt!*). There are *idioms*, like *Dave kicked the bucket* and *Kathy spilled the beans*. There are different kinds of *metaphors*, like *Susan flew down the street on her bicycle* and *That lecture was a sleeping pill*. There is *irony* and *sarcasm*, such as when your friend says *Now that was exciting*, after the sleeping pill lecture. All of these forms have in common the characteristic that what the speaker said, the *literal meaning*, is different than the interpretation that the speaker wants you to formulate, the *speaker meaning*. Because what is said differs from the intended meaning, nonliteral language requires the listener to draw *pragmatic inferences*, by combining information about the speaker, the context in which the expression is produced, and the literal content of the utterance to answer the questions: What meaning does the speaker wish to convey (what is on the speaker's mind)? Why did the speaker produce that utterance in this context?

How are these pragmatic inferences drawn so that the listener can see beyond the literal meaning to the speaker meaning? Answering this question entails solving two related problems. First, how do you know that the speaker does not intend a literal meaning? That is, how do you spot when the speaker is using nonliteral, rather than literal, language? This is sometimes called the *recognition* problem (Stern, 2000). Second, when the speaker meaning differs from the literal meaning, how does the listener compute the nonliteral meaning? Let's start with the recognition problem first.

The Standard Pragmatic View

Figurative language has traditionally been considered derivative from and more complex than ostensibly straightforward literal language.
SAM GLUCKSBERG

The *standard pragmatic view* assumes that computing literal meaning is the core function in language interpretation (Clark and Lucy, 1975; Glucksberg, 1998; Searle, 1979;

see also Holyoak and Stamenković, 2018). According to this view, when people hear a metaphoric expression, the first interpretation that the language comprehension system comes up with is the one that is most closely connected to tangible objects and the directly perceivable world. If someone says *Deb's a real tiger*, the standard pragmatic view says that you will interpret the expression as if it meant *Deb is a true example of a feline predator, usually orange-brown colored with black stripes, whose current range is limited to parts of south Asia and eastern Siberia*. Of course, normal people will *not* interpret *Deb's a real tiger* in this way—they will quickly come up with a more sensible interpretation. Nonetheless, the standard pragmatic view of nonliteral language interpretation argues that the literal meaning is computed and at least temporarily considered as the intended meaning of the expression, even though that initial interpretation is later discarded in favor of a more sensible one (e.g. *Deb has an energetic and fierce personality*).

The standard pragmatic view resembles the *garden path theory* of syntactic parsing. Both suggest an initial stage of interpretation that works for many expressions, followed by a second stage of interpretation involving re-analysis and re-interpretation for expressions that do not fit the normal pattern. According to the standard pragmatic view of nonliteral language interpretation, the reanalysis process involves computing an alternative meaning, rather than an alternative syntactic structure. The standard pragmatic view adopts this posture because, as Glucksberg explains, nonliteral meanings are viewed as being more mysterious and complex than literal meanings. Consequently, as suggested by Colston and Gibbs (2002, p. 58), "Under [the] standard pragmatic view, understanding any nonliteral utterance requires that listeners analyze a sentence's literal meaning before other figurative meanings can be derived."

Given the posture that literal meanings are computed first, how and why are metaphoric meanings ever discovered? Given that every expression can be interpreted literally (people really can kick the bucket, there really could be a Siberian tiger named "Deb" in your local zoo), how do people discover that the speaker or writer intended a nonliteral meaning?

According to the standard pragmatic view, solving the recognition problem involves computing the literal meaning first and then checking whether the literal meaning makes sense in context (Clark and Lucy, 1975; Miller, 1979; Searle, 1979; Stern, 2000). The listener first computes the literal meaning, and then attempts to integrate that meaning with the linguistic and social context. The listener will engage in further pragmatic inferencing and nonliteral interpretation only if the literal meaning is deficient in some way.

The question then becomes: How does the listener know whether the literal interpretation is deficient? One way would be to simply check whether the literal interpretation is true or false. If, as the Barenaked Ladies say, *You can be my Yoko Ono*, that must be literally false. There is only one Yoko Ono, and you are not her. If *You can be my Yoko Ono* is obviously false, then the listener will be motivated to answer the question: Why did the speaker say something that is obviously false? Sometimes, this conflict between the truth conditions imposed by the world and the information content of the utterance can be resolved by adopting a metaphoric interpretation. You are not literally Yoko Ono, but you can have some of the properties that Yoko Ono exemplifies. If you really do have some of the properties that Yoko Ono exemplifies, then the meaning assigned to that utterance is no longer false, and you have simultaneously discovered a true interpretation and the reason why the speaker said *You can be my Yoko Ono*. The speaker is expressing an opinion that you possess some of the qualities that Yoko Ono has, which could be true.

This solution of the recognition problem runs into trouble fairly quickly, however, because many literally true expressions are assigned a nonliteral meaning (Glucksberg

and Keysar, 1990; Glucksberg et al., 1992; Stern, 2000). When George Harrison of the Beatles says *I'm not the wreck of the Hesperus*, that is literally true. But Harrison is not rejecting a literal comparison of himself and a shipwreck, he is rejecting one metaphoric comparison in favor of a more apt metaphoric comparison (and so the next line of the song is *Feel more like the Wall of China*). Here's another example: *John is a real Marine* is literally true if John is in the Marine Corps. Nonetheless, most people would adopt a nonliteral interpretation of that expression, along the lines of *John is a really good, exemplary, or skilled Marine*. Similarly, the expression *My wife is an animal* is literally true. But a nonliteral meaning, along the lines of *My wife behaves in an unpredictable and uncivilized way*, will be preferred over the literal interpretation. Literal falsehood is not a necessary precondition for an utterance to be assigned a nonliteral meaning. What we need, then, is a theory that does not require utterances to be literally false before they are assigned a nonliteral interpretation.

If literal falsehood is not sufficient to classify an utterance as nonliteral, how can we know when a nonliteral meaning is intended? There are, in fact, a number of other criteria that could be applied. As Stern (2000, p. 3) explains, an utterance could be considered deficient if its literal meaning is "grammatically deviant, semantically anomalous, explicitly or implicitly self-contradictory, conceptually absurd, nonsensical ... pragmatically inappropriate, obviously false, or so obviously true that no one would have reason to utter [it]." According to the standard pragmatic view, the listener will consider a nonliteral interpretation when the literal meaning lacks one or more of the characteristics that make utterances fit into their contexts. For instance, if you were at a party and an attractive stranger said *Can you get me a beer?* you wouldn't just say *Yes* (which is the appropriate response to a literal request for information). You would go and fetch a beer. On the other hand, if the same attractive stranger said *Can you bench press 200 pounds?*, you might very well say *Yes I can, by cracky! Come by the gym tomorrow and I'll show you.*

According to the standard pragmatic view, the indirect question *Can you get me a beer?* violates the normal conversational rules (as laid out by philosophers like H. Paul Grice, 1989) because people at parties generally obviously have the physical ability to fetch beer. The indirect request is therefore deviant, because the answer to the question (on its literal interpretation) is so obvious that no sane person would ever ask it. The listener can repair the situation, however, by reinterpreting the "deviant" utterance as a (nonliteral) indirect request. By contrast, most people at parties do not obviously have the physical ability to lift very heavy weights, and so *Can you bench press 200 pounds?*, which has the same form as *Can you get me a beer?*, is not considered deviant, and is therefore interpreted as a literal request for information.

A literal interpretation of an utterance may be rejected in favor of a nonliteral meaning if the literal interpretation is false (as in *Deb is a real tiger*) or if the literal interpretation violates one or more of the characteristics of "normal" utterances. What are the characteristics of "normal" utterances, and where do they come from?

One general theory that explains why utterances are "normal" or "abnormal" is Grice's (1989) theory of *conversational maxims*. A conversational maxim is a rule or guideline that applies to things people say while having a conversation. Grice starts with the idea that participants in a conversation try to cooperate with one another in order to expand the pool of shared knowledge (see also Clark, 1996; Glucksberg and Keysar, 1990; Wilkes-Gibbs and Clark, 1992). Engaging in this cooperative activity involves following other rules (maxims) as well. These include the maxims of *quality*, *quantity*, *manner*, and *relation*. The maxim of *quality* specifies that you should tell the truth; your utterances should be literally true. We don't like it when people lie to us.

The maxim of *quantity* specifies that your utterances should provide new information. So you should not simply repeat information that is already in *common ground* (knowledge that is shared, and known to be shared, between the participants in the conversation). We get annoyed when we have to listen to the same story over and over again, or when someone repeats themselves unnecessarily, or belabors a point long after we have already figured out the thrust of their argument, or tries our patience by sticking to a settled topic, or doesn't tell us anything that we don't already know.

The maxim of *manner* specifies that your utterances should be clear and unambiguous. You should convey information as plainly and directly as possible, so that your utterance does not have multiple possible interpretations.

The maxim of *relation* specifies that your utterance should contribute to or continue the current topic of discussion unless you explicitly introduce a new topic. So, if we are talking about baseball, my next utterance should be on the topic of baseball unless I say something like *Enough baseball, let's talk about me now.*

Indirect requests such as *Can you get me a beer?* can violate the maxim of quantity. It is obvious from my appearance that I am physically capable of retrieving beer, so uttering that question does not provide me with an obvious means of moving the conversation forward *if* it is interpreted as a literal request for information. Metaphoric expressions also commonly violate one or more of the Gricean conversational maxims. If someone says to you *My wife is an animal*, the literal interpretation will violate at least the maxims of quantity and relation, and possibly quality and manner. If you intended to use the literal meaning of *animal*, saying that your wife is an animal does not provide any new information (by definition, a wife must be a female human, and humans are a type of animal), so the utterance violates the maxim of quantity (add *new information*). It would be like saying *My wife has lungs and skin*, or *My wife has eyelids.* Likewise, if the topic of conversation is one's spouse or one's relatives, it is hard to see how the information that someone is an animal relates to the general topic. Therefore the utterance violates the maxim of relation (*stick to the topic*). So, while literal falsehood may not be a sufficient test for nonliteral intentions, Gricean violation could be. If so, solving the recognition problem involves (a) computing the literal meaning of a given expression and (b) checking that literal meaning against the requirements imposed by Gricean maxims.

Whether the appropriate standard is literal falsehood or Gricean violation, the standard pragmatic view argues that we attempt a nonliteral interpretation only after first computing and testing a literal interpretation against the preceding context. This solution to the recognition problem, then, naturally leads to a set of processing assumptions under which literal meaning is the default, and nonliteral meaning is optional. One prediction that follows is that people, when listening to nonliteral language, will only compute the intended nonliteral meaning after they compute and reject the (unintended) literal meaning. But is it true that literal meanings are always computed first?

A substantial body of experimental evidence shows that nonliteral meanings are computed as fast as literal meanings. Ray Gibbs assessed the interpretation of indirect requests, such as *Can you open the window?* (Gibbs, 1983), to see whether the direct, literal meaning was computed before the nonliteral meaning. In a reaction-time experimental task, participants read either literal (e.g. *I would like you to open the window*), direct requests or nonliteral, indirect requests (e.g. *Can you open the window?*). The participants' task was to judge, as quickly as possible, whether paraphrases expressed the same meaning as the direct or indirect requests. Gibbs measured how long it took people to judge the paraphrases, on the assumption that this judgment could be done only after an interpretation of the direct and indirect requests had been computed. If literal

meanings are computed before nonliteral meanings, then paraphrase judgment times should be shorter in the direct request condition than in the indirect request condition. (Because the paraphrases were the same across the literal and nonliteral conditions, the amount of time it took the subjects to interpret the paraphrases should not have influenced the outcome of the experiment.) If nonliteral meanings are computed only when literal meaning is computed and found deficient, then paraphrase judgment should take longer in the indirect request condition than in the direct request condition. Gibbs found that paraphrase judgment times were the same across conditions, showing that his subjects interpreted the indirect requests just as quickly as the direct requests. In a related study, participants were asked to paraphrase literal and metaphoric expressions (Harris, 1976). The amount of time it took participants to start paraphrasing (*paraphrase initiation time*) was measured, on the assumption that paraphrases could be initiated only after a meaning was computed. Paraphrase initiation times were the same for literal and metaphoric expressions, indicating that both literal and nonliteral expressions were computed at about the same speed. In another experiment using a sentence categorization task, subjects were able to identify and classify metaphoric expressions as quickly as they identified literal statements (H. Pollio et al., 1984).

Data from other experiments also support the idea that nonliteral meanings can be computed as quickly as literal meanings. In the Blasko and Connine (1993) study, participants were presented with novel metaphoric expressions, such as *indecision is a whirlpool*. The question was whether the literal meaning of the final word *whirlpool* would be accessed before a nonliteral or metaphoric meaning. The literal meaning of *whirlpool* is a mass of water that is circling and creating a depression in the surface of a body of water. To see whether people were thinking of that literal meaning, participants heard the phrase *indecision is a whirlpool* and responded to the visually presented target word *water*. In the context of *indecision is a whirlpool*, the metaphoric meaning of *whirlpool* doesn't really have anything to do with water. Instead, the metaphoric meaning is something like *When people are indecisive, their thinking goes around in a circle, and they act confused*. To see whether people were thinking of that nonliteral meaning, subjects responded to the target word *confusion* after hearing the phrase *indecision is a whirlpool*. If literal meanings are computed faster than nonliteral meanings, then subjects should respond to the literal-related target word *water* faster than the nonliteral-related target word *confusion*. In fact, people responded to target words related to nonliteral meanings (e.g. *confusion*) just as fast as they responded to target words related to literal meanings (e.g. *water*); and both kinds of target words were responded to faster than control words that were totally unrelated to any meaning of the expression *indecision is a whirlpool*. These results therefore support the idea that nonliteral meanings are computed just as quickly as literal meanings.

McElree and Nordlie (1999) applied an unusual but highly effective experimental technique, speed-accuracy tradeoff (SAT), to find out how quickly literal and metaphoric meanings were computed (see also Bambini et al., 2021). SAT provides a very accurate way of measuring when, exactly, different kinds of information become available and start to influence a person's behavior (see the box on p. 277 for a summary of how it works). McElree and Nordlie had people read literal and nonliteral expressions and used SAT to measure how quickly people could understand the different kinds of expressions. The SAT results for literal and nonliteral expressions were virtually identical, showing that nonliteral meanings were computed just as quickly as literal meanings.

The previously summarized studies might all be viewed as involving unusual circumstances, and so you might want to discount their support for simultaneous computation of literal and nonliteral meaning. Normally, when we encounter expressions in language,

we are not asked to judge paraphrases, respond to target words that pop up separately from the main conversation, or respond on somebody else's deadline. Fortunately, there are more naturalistic methods that point toward the same conclusions.

Consider a sentence-by-sentence reading study by Ortony (1979). In this study, participants read target sentences like (1).

(1) The sheep followed their leader over the cliff.

Sentence (1) can be interpreted as referring to real flesh-and-blood farm animals and a real geological feature, but it can also be interpreted in nonliteral terms. If sentence (1) followed sentence (2), it should be assigned a nonliteral interpretation.

(2) The investors looked to the Wall Street banker for advice.
(1) The sheep followed their leader over the cliff.

In the context of sentence (2), *sheep* refers to *investors*, *their leader* refers to *the Wall Street banker*, and *the cliff* refers to a sharp drop in the value of investments. So, to inter-

SPEED–ACCURACY TRADEOFF

In most reaction-time experiments participants respond at their own speed. They are often encouraged to respond as quickly as possible without making errors, but it is still up to each individual subject to decide how to weigh these two criteria. Some subjects respond very quickly, before they are really ready, and they make a lot of errors. Data from that kind of subject is often discarded before the data from a study are analyzed. Some subjects respond more slowly, only after they are really, really sure that they have the right response ready, and they usually make very few errors. Many subjects with this profile are commonly included in published studies. If too many such subjects wind up in the analyzed data, the results may overestimate the amount of time it takes to complete an information processing task. As a result, one problem in interpreting many reaction-time studies is that we do not know to what extent participants were trading extra time for increased accuracy, or decreased accuracy for greater speed. The SAT paradigm solves this problem. It does so by preventing subjects from trading speed for accuracy. Instead, subjects are trained to respond before a set, fixed deadline. If the deadline is very long, subjects' accuracy is very high. If the deadline is very short, subjects' accuracy will be very low. At intermediate deadlines, subjects' accuracy will be greater than zero, but less than perfect. By testing a set of subjects on many hundreds—sometimes thousands—of trials with different deadlines, experimenters can see when subjects' accuracy diverges from chance performance. At the point in time where subjects' performance becomes better than chance, the source of information that was manipulated must have been available to influence the subjects' behavior before that point (see Martin and McElree, 2018; McElree and Griffith, 1995, 1998; McElree and Nordlie, 1999; McElree, Pylkkänen et al., 2006).

pret sentence (1) in the context of sentence (2), people have to assign nonliteral meanings to a number of the words in sentence (1). That condition was contrasted with a condition in which the literal meanings of the same words were called for, which is what happened when sentence (1) followed sentence (3).

(3) The animals were grazing on the hillside.
(1) The sheep followed their leader over the cliff.

Ortony measured how long it took people to read sentence (1) when it followed sentence (2), the *nonliteral* condition, and when sentence (1) followed sentence (3), the *literal* condition. If literal meanings are computed faster than nonliteral meanings, then reading times for sentence (1) should have been shorter in the literal condition than in the nonliteral condition. In fact, subjects in this experiment read sentences like (1) just as quickly in the nonliteral condition as in the literal condition. These results offer further evidence that nonliteral meanings are computed just as quickly as literal meanings.

Other eye-tracking studies looked at processing of individual words or phrases within sentences, rather than whole-sentence reading time (Inhoff, 1984; Shinjo and Myers, 1987), and also found equal reading times for components with literal meanings and components with nonliteral meanings. Because reading is a highly practiced and very natural task for most adults (and especially the college students who served as subjects in these experiments), these studies are not subject to the "naturalness" criticism that might be applied to some of the other investigations of nonliteral language processing, and they reach the same conclusions. On balance, research on nonliteral language processing shows that the "literal first" assumption of the standard pragmatic view cannot be correct.

The standard pragmatic view also views nonliteral interpretation as being optional in contexts where the literal interpretation fits reasonably well with the preceding context. However, research using a semantic version of the Stroop task shows that nonliteral interpretation is obligatory, not optional, even in contexts where the literal interpretation would suit the comprehender's needs (Glucksberg, 1998, 2003; Glucksberg et al., 1982; Keysar, 1989; Stroop, 1935; see also Kazmerski et al., 2003; Wolff and Gentner, 2000). In the Stroop task, color names are printed in ink that does not match the color name. So, the word *blue* is printed in red ink. The subject's task is to say the name of the color of the ink, and ignore what the word says. So, you look at the word *blue* printed in red ink, and you have to say red, rather than *blue*. Subjects experience great difficulty ignoring the color name *blue*, so they are slow to say the ink color *red*, and they tend to make errors. This task shows that access to the color name *blue* occurs automatically when people look at the word *blue*.

Are nonliteral meanings also automatically computed or accessed? To find out, people read statements that were literally false, but had a good nonliteral interpretation. For example, in a context where *Keith* is described as an adult who acts in an immature way, the statement *Keith is a baby* is literally false, but has a good nonliteral interpretation. People who read statements such as this had the task of identifying statements that were literally true. In that task, the correct response to the statement *Keith is a baby* is to say "false." This task resembles the Stroop task, in that automatic access to the good metaphorical meaning could make the subjects' job harder, because the good metaphorical interpretation makes it harder for them to say that the sentence is untrue. Even though the task is to look for literal meaning only, if the metaphoric meaning is plausible, and if the metaphoric meaning is computed automatically, then people will have a hard time rejecting *Keith is a baby*, the same way they have trouble rejecting the response *blue*

when the word *blue* is printed in red ink. In fact, people had much greater difficulty rejecting metaphorically "true" statements like *Keith is a baby* than equivalent literally and metaphorically "false" statements, such as *Keith is a basin*. So, in the context of a task where literal meaning was sufficient (people could do the judgment task without ever computing nonliteral meaning), people, nonetheless, automatically computed the nonliteral meaning.

The preceding experiments show that the standard pragmatic view has two major things wrong with it. First, nonliteral meanings become available to the listener as quickly as literal meanings do.[2] The assumption that literal meanings are computed first is not supported by the experimental data. Second, the standard pragmatic view says that computation of nonliteral meanings is optional and undertaken only when the literal meaning is problematic in a given context. However, experiments like Glucksberg and colleagues' (Glucksberg, 1998, 2003; Glucksberg et al., 1982) show that nonliteral meanings are computed automatically in contexts where the literal meaning is entirely sufficient for the task at hand. Results like these have persuaded many language scientists to adopt an alternative theory under which literal and nonliteral meanings are computed in parallel (e.g. Swinney and Cutler, 1979). Simultaneous computation of literal and nonliteral meaning is consistent with the experimental results showing that both kinds of meanings are computed equally quickly and that nonliteral meaning is computed whether the literal meaning is true or false. So, interpreting nonliteral meanings is like interpreting words with multiple meanings. In both cases, multiple meanings are automatically activated and subsequently assessed against the requirements imposed by context, with the most compatible meaning eventually being selected as *the* interpretation of the ambiguous expression. We will explore the multiple meanings associated with nonliteral expressions and how the system might go about selecting one contextually appropriate meaning after we discuss a bit more about one specific kind of nonliteral expression: metaphor.

Metaphor

A metaphor allows us to indicate a relationship between two elements. The first element is the *topic* and the second element is the *vehicle*. The topic is the focus of the conversation; it is what the discourse is about (Ortony, 1975). The vehicle is some concept or exemplar that we are using to comment on the topic. In the expression *Nicole Kidman is bad medicine, Nicole Kidman* is the topic and *bad medicine* is the vehicle. Of course, the expression is *not* intended to be interpreted as saying that Nicole Kidman is a pharmaceutical product that comes in bottles and has harmful side effects. Instead, the nonliteral speaker meaning is something like *Nicole Kidman is dangerous and can make you feel really bad.*[3]

This expression is in the commonly occurring *A is B* form frequently used to express metaphoric meanings. Note that metaphors of this type are subject to the recognition problem because literal category-inclusion statements also take the form *A is B*, as in *Copper is a metal*, and *Dogs are mammals*.

How do we solve the recognition problem for metaphors? This is not an easy question to answer in the abstract, but when we have a metaphoric comparison involving topic *A* and vehicle *B*, *A* and *B* are in some unusual relation to one another. As Greg Murphy (1996, p. 175) explains, in reference to the metaphoric expression *Lee is a block of ice*:

> *In order for a sentence to be perceived as metaphoric, the vehicle cannot apply in a straightforward way to the topic. For example,* Lee is an attorney *does not … require any special ground for its interpretation. The usual, familiar meaning of* attorney *specifies a person with a particular profession, and since* Lee *is the name of a person, there is no inconsistency in calling Lee an attorney. However,* block of ice *literally means the solid, frozen state of the substance H_2O, and since a person is typically neither H_2O nor frozen, this predicate cannot be applied to Lee in a straightforward way. There must be some kind of mapping from the usual meaning of this phrase to the conveyed, nonliteral meaning.*

The *A is B* form is one kind of metaphor, called an *attributive metaphor* (Glucksberg and Keysar, 1990). When you use that kind of expression, you are asserting that some attributes (properties, characteristics, or features) of the vehicle apply to the topic. Sometimes, we use *single-word metaphors*, as in the expression *The girl flew down the hill on her bicycle.* In this case, the verb *flew* is not used in its default sense, as indicating motion through the air, but rather in a metaphoric sense that implies very rapid motion on the ground. We can also use *relational metaphors*, which you are familiar with from your experience with standardized tests like the ACT and the SAT. Relational metaphors take the form: *A is to B as C is to D* (or the shorthand *A:B::C:D*). If you picked up an article about the war in Ukraine, you might read, *President Biden considered Vladimir Putin a modern-day Hitler.* This is actually a nifty rhetorical device because it sets up an implicit relationship between Mr. Biden and an unmentioned third person (besides Putin). See if you can work out what that relationship might be and who the unmentioned third person is (see endnote 4 for the spoiler).[4]

How do we figure out what meaning we should apply to a metaphor? Do we have special semantic interpretation processes that we use for metaphors but not for literal expressions? According to the *comparison* school of thought (e.g. Bottini et al., 1994; Fogelin, 1988; Ortony, 1979; Tversky, 1977), people interpret metaphors by mentally converting them to similes. Similes are like metaphors in that they both have a topic and a vehicle; but similes are unlike metaphors in that they have an extra word (*like*, usually), and that extra word makes it explicit that two things are being compared. So, similes take the form *A is like B* (or *A resembles B*, or *A is similar to B*). We can use similes to point out that two things have shared characteristics. You could say, *Copper is like tin, Baseball is like cricket*, or *Mexico is like Spain*. These are considered literal comparisons, because the default meanings of the topic and vehicle are involved in the comparison. Copper is like tin, because the default meanings for both include the information that they are metals, fairly common, dug out of the ground, only a little bit shiny, and so forth. Baseball and cricket both involve teams of players, balls, bats, grass fields, and so on.

When we use a metaphor, such as *Nicole Kidman is bad medicine*, perhaps we interpret that expression by mentally converting it to a simile, of the form *Nicole Kidman is like bad medicine*. If that were the case, then after mentally converting the metaphor to a simile, we could apply the same processes to evaluate the metaphoric relationship between Nicole Kidman and bad medicine as we use to evaluate the literal relationship between copper and tin. One advantage of this approach is that we need only a bare minimum of special-purpose interpretive machinery for metaphors—the only special-purpose mental process is the one that turns metaphors into similes, and there the conversion is as simple as changing *is* to *is like*. In fact, this last point bears repeating. The comparison approach to metaphor comprehension asserts that the meaning of a metaphor is discovered by converting it to a simile, and the meaning stays the same before and after the conversion—the two expressions convey the same meaning. The only thing that changes between a metaphor and a simile is how the expression is worded.

Although the conversion-to-simile view of metaphor interpretation has the advantage of being simple and straightforward, it may not be able to account for the full range

of facts about similes and metaphors. In fact, a closer examination of literal comparisons and metaphoric comparisons quickly turns up substantial differences between the two. One difference is that literal comparisons are generally *reversible*, but metaphoric comparisons are not. In a literal comparison, most of the time, it does not matter which element appears in which order; the basis of comparison between the two elements stays the same, and the overall semantic force of the statement also remains the same. *Copper is like tin* is essentially equivalent to *Tin is like copper*. There is no strong basis for preferring one version over the other, except perhaps if one of the two elements is made topical. So you might prefer to say *Baseball is like cricket*, if your English friend asked you to explain the rules of baseball. But you may prefer the opposite, *Cricket is like baseball*, if cricket were the topic of conversation.

Many metaphors are not reversible in this way. You can say *Nicole Kidman is bad medicine*, but you can't say *Bad medicine is Nicole Kidman*. Some metaphors are reversible, but the meaning of the expression changes when it is reversed (unlike the similes, where the meaning of the expression as a whole remains essentially unchanged). For instance, if you said My *surgeon is a butcher*, the *grounds of comparison*, the relationship that connects *surgeon* to *butcher*, is that the surgeon is a bad surgeon (he chops up meat willy-nilly, like a butcher). If you reverse the topic and the vehicle, you get *My butcher is a surgeon*, which implies that the butcher is highly skilled. The same thing applies if these metaphors were converted to similes (*My surgeon is like a butcher; My butcher is like a surgeon*). So, there is something fundamentally different in literal comparisons and metaphors. Literal comparisons have more stable grounds for comparison and so the elements can be reversed; metaphors call on different grounds for comparison depending on the precise nature of the topic and vehicle and are therefore not generally reversible (see also Zhou and Tse, 2022).

The conversion-to-simile view of metaphor interpretation also predicts that metaphoric expressions will take longer to interpret than similes that express the same relationship between topic and vehicle (because there is an extra conversion step necessary to comprehend metaphors). However, reaction-time studies have shown that, under some circumstances, similes take *longer* to understand than equivalent metaphors (Glucksberg, 1998, 2003). It appears therefore that metaphors can be interpreted without mentally converting them to similes.

To interpret metaphoric expressions, the listener must discover the grounds of comparison that the speaker is using to connect the topic to the vehicle. How do listeners accomplish this feat? The *property-matching hypothesis* explains why and how similes and metaphors communicate relationships between topics and vehicles (Johnson and Malgady, 1979; Malgady and Johnson, 1980; Miller, 1979; Ortony, 1979; Tversky, 1977). According to the property-matching hypothesis, interpreting similes and metaphors depends on finding properties of the topic that are identical to properties that the vehicle has. One advantage of this approach is that property matching can be used to interpret both literal comparisons and metaphoric ones, so it does away with the need for any special interpretation processes that apply only to metaphors. To interpret literal comparisons, such as *a dog is a mammal*, you find properties that you know go with the vehicle *mammal*, such as *has fur, has mammary glands, bears its young live*, and you see whether there are matching properties listed under the entry dog. When you find matching properties, you highlight them to draw attention to the fact that they are shared between the topic and vehicle. You do the same thing for metaphors. When you hear *Nicole Kidman is bad medicine*, you find properties of the vehicle (*dangerous, makes people feel bad*) and you look for the same properties under the entry *Nicole Kidman*. The same process can be used whether the comparison is based on the metaphor form, *A is B*, or the simile form, *A is like B*. The property-matching hypothesis also can explain why some metaphors just don't make sense. If no shared properties can be found for the topic and vehicle,

then the metaphor will not make sense. If someone said, *billboards are like pears* (or *billboards are pears*), that would sound strange. According to the property-matching hypothesis, *billboards are like pears* sounds strange because there are no properties that *pears* and *billboards* have in common, and so there are no grounds for comparison.

The *salience imbalance* hypothesis (Johnson and Malgady, 1979; Tourangeau and Sternberg, 1981) is a refined version of the property-matching hypothesis. This hypothesis attempts to explain why speakers choose the metaphor (*A is B*) form sometimes and the simile (*A is like B*) form other times when highlighting relationships between topics and vehicles. Specifically, the salience imbalance hypothesis proposes that literal comparisons, of the form *A is like B*, will be used when the grounds of comparison involve properties that are salient in both the topic and the vehicle. Metaphors, of the *A is B* form, are used when the topic and vehicle share properties, but the grounds of comparison involve properties that are obscure in the topic, but highly salient in the vehicle. According to the salience imbalance hypothesis, we use *copper is like tin*, because being a metal, being mined, being useful in manufactured products, and being only a little bit shiny are salient properties of both copper and tin. That is, if we asked people to list properties of copper and tin, *metal, mined*, and *sort of shiny*, would appear high on the list for both copper and tin.

The topic and vehicle in a metaphor, by contrast, produce different lists of properties, with different characteristics appearing in very different orders for the topic and vehicle. If, for instance, you were asked to list the properties of *Nicole Kidman*, you would most likely include *actress, dancer, Australian, cute and sassy, performed brilliantly opposite Will Ferrell in "Bewitched"* near the top of the list and *heart-breaker, makes people feel bad* somewhere down near the bottom. Suppose you were then asked to list the properties of *bad medicine*. Right at the top would be properties such as *dangerous* and *makes people feel bad*. So, *makes people feel bad* is a highly salient property of *bad medicine*, where salience is operationalized as position on the list. (Properties that are listed early are salient because, presumably, they come to mind easily.) When someone says, *Nicole Kidman is bad medicine*, *bad medicine* has salient properties (*dangerous, comes in bottles, makes people feel bad*) that match non-salient properties of Nicole Kidman (and *comes in bottles* can be ruled out as part of the grounds of comparison, because no such property will be found under the *Nicole Kidman* entry). By putting *Nicole Kidman* and *bad medicine* together, the speaker is promoting one or more low-salience features of Nicole Kidman into much higher and more salient positions in the mental representation of the topic. In this way, metaphoric expressions make the listener think about the familiar in a new way, by reordering the salience of different properties associated with the topic, which contributes to the rhetorical force of the metaphoric expression.

Although property matching, like the conversion-to-simile view has the advantage of explaining how literal and metaphoric comparisons are interpreted, and why speakers would prefer a simile in one instance and a metaphor in another, the property-matching hypothesis has trouble explaining the full range of metaphoric expressions. One problem is that the property-matching hypothesis predicts that that the grounds of comparison in metaphors should involve low-salience properties of the topic and high-salience properties of the vehicle, but sometimes metaphors that are easy to interpret involve properties that are *low* in salience in both the topic and the vehicle. Consider sentence (4) (from Glucksberg and Keysar, 1990):

(4) The senator was an old fox who could outwit the reporters every time.

What is the ground of comparison here? It must be that the senator is *clever* or *wily*. If you asked people to list properties of the word *senator*, you would get a list that included *politician, powerful, distinguished*, and so on. If you asked people to list properties of *fox*,

you would get a list that included *furry, bushy tailed, hunted with dogs, can't be trusted to guard the hen-house*, and so on. *Clever* would be found way down near the bottom of properties listed in response to both *senator* and *fox*, if that property appeared at all. According to the salience imbalance hypothesis, sentence (4) should be a lousy metaphor, but people actually rate that sentence as being a good or *apt* metaphor. So an imbalance between the salience of a property in the topic and vehicle is not a necessary precondition for a good metaphor.

Even worse for the property-matching view, some really good metaphors involve topics and vehicles that have zero shared properties. One standard metaphor is *No man is an island*. You could ask people to list the properties of *man* and *island* as long as you like, and it is unlikely that you will find any shared properties besides the fact that both *man* and *island* are *things* or *nouns* (and neither of those is involved in the grounds of comparison for *No man is an island*). Metaphors of the form *no man is an island* illustrate the principle that metaphors can be used to assign brand new properties to the topic, rather than merely highlighting existing, low-salience properties. Currently, the property-matching hypothesis does not have any way to accommodate this phenomenon, even though *attribute introducing metaphors*, such as *the mind is a computer*, have played central roles in both common and scientific advances in knowledge. As Bowdle and Gentner (2005, p. 194), explain, "The computer metaphor of mind was not informative because it simply highlighted certain well-known aspects of the mind that were also true of computers but rather because it promoted a transfer of knowledge from the domain of computers to that of minds. Feature-matching models provide no mechanism for the projection of such distinctive (vehicle) properties."

A further problem with the comparison view is that metaphor and simile versions of a comparison do not always convey the same meaning because meaning does not always survive the conversion process. Literal similes cannot be translated into class-inclusion statements at all. You can't convert *copper is like tin* into *copper is tin* (Glucksberg and Keysar, 1990). Likewise, many metaphors cannot be converted to similes and retain their meaning. Consider the expression *My lawyer is a well-paid shark* (Bowdle and Gentner, 2005; Glucksberg and Haught, 2006a, 2006b). People assign expressions like that high aptness ratings (they think it is a good metaphor). But that expression does not fare so well when it is converted to a simile. People assign much lower aptness ratings to the simile version *My lawyer is like a well-paid shark* than they assign to the metaphor version. So, metaphor and simile forms do not have equivalent meanings, and that goes against the comparison and mental-conversion-to-simile views of metaphor interpretation. As Glucksberg and Haught explain (Glucksberg and Haught, 2006a, p. 376), "If metaphors cannot always be paraphrased as similes, then metaphors cannot, in principle, be understood in terms of their corresponding similes, and vice versa. This means that comparison theories of metaphor comprehension, which rest on the assumption that metaphors and similes are equivalent, are fundamentally flawed."

Class inclusion and dual reference

The *class-inclusion* hypothesis of metaphor interpretation provides an alternative to the conversion-to-simile and property matching views (Glucksberg, 1998, 2003; Glucksberg and Haught, 2006a, 2006b; Glucksberg and Keysar, 1990; Glucksberg, McGlone et al., 1997). According to the class-inclusion view of metaphor interpretation, all expressions, literal and nonliteral, of the form *A is B* are interpreted as asserting that the topic (*A*) is a member of the category represented by the vehicle (*B*). This works equally well for both

literal and metaphoric comparisons. A speaker can say *A dog is a mammal*, and you can understand that by recognizing that *dog* is a member of the category *mammal*, and therefore inherits the characteristics of members of that category. Nonliteral metaphoric expressions can be interpreted in a similar fashion. When I say, *Nicole Kidman is bad medicine, bad medicine* is a prototype for the ad hoc category *things that are surprisingly dangerous and bad for you*. If someone says, *Kyle is a pop-up ad*, you can interpret that by placing *Kyle* in the category that *pop-up ad* exemplifies—unexpected, mildly annoying things that temporarily stop you from doing whatever it is you're doing. The class-inclusion hypothesis has the advantage of providing a unified explanation for how metaphoric and literal statements are interpreted, and it avoids some of the pitfalls of the comparison and mental-conversion-to-simile approaches.

The class-inclusion view also has the advantage of explaining how and why people respond to metaphors and similes, and why there is a difference between metaphors and similes. To do so, it requires one auxiliary assumption. Namely, metaphoric expressions involve *dual reference* (a bit like ambiguous words), while literally interpreted similes and class-inclusion statements involve only a single mapping between the topic and the vehicle. According to the dual reference hypothesis, when you see a statement of the form *My lawyer is a shark*, you think of a literal, flesh and blood, cartilaginous fish and you simultaneously think about an ad hoc category of dangerous and aggressive animals (which could also include, lions, tigers, and bears; as in *My lawyer is a lion/tiger/bear*). *Shark*, in this case, refers to both a *basic-level* concept (a real shark) and a *superordinate category* that the shark exemplifies, or is a prototypical member of. The simile version, *The lawyer is like a shark*, makes reference to just the basic level, literal *shark* and *not* to the superordinate category. So, unlike the metaphoric expression, the simile version invites the question *The lawyer is like a shark in what way?*

The fact that metaphoric expressions involve dual reference and similes involve single reference means that you can do some things with metaphoric expressions that you cannot do with similes. Because *My lawyer is a shark* simultaneously refers to a real shark and a category of dangerous things, you can use a modifier, like *well-paid* to modify shark, as in (5) (Glucksberg and Haught, 2006a, p. 368):

(5) My lawyer is a well-paid shark.
(6) *My lawyer is like a well-paid shark.

The simile version in (6), by comparison, seems a bit odd. This is because in the simile version, *shark* refers to the literal shark, and its reference does not extend to the superordinate ad hoc category *dangerous things*. Since real, live, swimming sharks don't draw paychecks, it does not make sense to say *well-paid (literal) shark*, and the expression as a whole seems weird. This sense of weirdness was confirmed in an experiment where subjects were asked to rate the quality of expressions such as (5) and (6). Subjects rated modified metaphoric expressions like (5) higher than modified simile expressions like (6). Subjects also read expressions like (5) faster than expressions like (6), even when adjustments were made to account for the extra word in (6).

Further evidence that metaphoric expressions of the form *A is B* are interpreted as class-inclusion statements that call to mind superordinate categories comes from priming experiments (Glucksberg, Manfredi et al., 1997). Participants' reading times were measured for metaphoric expressions such as *My lawyer is a shark*. That target sentence could be preceded by a neutral control sentence or by a sentence that focused participants' attention on the literal meaning of the word *shark*. An example of a literal-focusing prime sentence would be *Sharks can swim*. The prime sentence *Sharks can swim* draws participants' attention toward the literal meaning of shark, and away from the superordinate

category *dangerous animals*. Under those conditions, participants had a harder time understanding the metaphoric expression *My lawyer is a shark*, which requires them to establish a connection between the topic (*lawyer*) and the superordinate category (*dangerous animals*).

The basic version of the category-inclusion hypothesis explains how a relationship is established between metaphorical topics and vehicles, but it does not explain how the same vehicle can point to different superordinate categories for different topics. For example, as Bowdle and Gentner (2005) point out, the vehicle *snowflake* can point to one superordinate category when the topic is *a child* and another category when the topic is *youth*. In the metaphoric expression *a child is a snowflake*, the metaphoric interpretation is that *every child is unique* (no two snowflakes are identical). In the metaphoric expression, *youth is a snowflake*, the metaphoric interpretation is that *youth is fleeting* (snowflakes melt easily).

To explain how the same vehicle can induce different superordinate categories, Glucksberg and his colleagues (Glucksberg, Manfredi et al., 1997) developed a version of the category-inclusion hypothesis under which the vehicle makes a set of superordinate categories available for interpretation, and characteristics of the topic point the reader toward the appropriate one. In that case, the topic and vehicle both impose constraints on interpretation, as opposed to the vehicle being the sole factor that influences the choice of superordinate category information. Because expressions such as *a child is a snowflake* and *youth is a snowflake* are equally easy to understand, Glucksberg and colleagues proposed that vehicles activate multiple superordinate categories in parallel. If superordinate categories were accessed one at a time, one of those metaphors should be easier to understand than the other.

The class-inclusion hypothesis of metaphor interpretation closely aligns the processing and interpretation of metaphoric expressions with literal expressions. Glucksberg therefore joins a growing group of language scientists who view metaphor as a normal, mainstream part of human language use, rather than being a peculiar class that needs special interpretive machinery and processes. In fact, rather than being rare and unusual, metaphors are produced frequently in the course of normal discourse. Speakers produce about six metaphors (about four "frozen" and two "novel" metaphors) per minute of speaking time, or about one every 10 seconds (H. Pollio et al., 1977).

Conceptual mapping and meaning

The standard pragmatic view proposes that metaphoric expressions are special and different from "normal," common literal expressions. The experimental evidence reviewed suggests that a number of assumptions that follow from the standard pragmatic view may not be correct. In light of this experimental evidence, a number of philosophers and language scientists have attempted to formulate new ideas about how metaphoric language fits into language use as a whole. One idea that has gained popularity is the *conceptual-mapping* hypothesis (Gibbs, 1994; Lakoff and Johnson, 1980a, 1980b). First, the conceptual-mapping hypothesis minimizes the distinction between metaphoric and literal language. Hence, philosophers, linguists, and psychologists who subscribe to the conceptual-mapping hypothesis view metaphors as being a ubiquitous feature of normal discourse, and not particularly special (e.g. Giora, 2007; Lakoff, 1987; Pinker, 1994). According to this view, much of the language that we view as literal is really based on implicit metaphors (as in the stoplight example at the beginning of this chapter). Second, the conceptual-mapping view suggests that much of our understanding of normal,

everyday words is based on discovering or highlighting links between different domains. Normal, literal words like *argument, love*, and *anger* are understood because they bring to mind other concepts, such as *war, journey*, and *heated fluid in a container*. This connection of different domains when we think about a concept such as *argument* allows us to use expressions like *his criticism was right on target, she attacked every one of my strong points*, and *my barely adequate psychic defenses crumbled in the face of her logical blitzkrieg*. According to the conceptual-mapping view, we can talk about arguments as war because the way we think about arguments, the way they are mentally defined, is based on connecting aspects of the process of arguing to analogous aspects of the process of fighting a war.

The conceptual-mapping hypothesis comes in *strong* and *weak* forms. According to the strong form, unless a word refers to a directly perceivable concept (like red), the word will be defined and understood because of its metaphoric relationship to some other, more basic domain. As Murphy (1996, p. 178) explains, "one does not really understand an argument—one only understands war, and the understanding of arguments is parasitic on this concept." In particular, certain *fundamental metaphors* are the basis of many commonly used expressions. These fundamental metaphors include *space and movement*, as well as *force, agency, and causation*. According to a *strong* form of the *conceptual mapping hypothesis*, all utterances other than those that directly express fundamental metaphors are interpreted by activating a fundamental metaphor and drawing connections (mapping) between the more complex domain and the fundamental metaphor (the *source domain*).

When you interpret the utterance *The meeting went from 3 to 4 o'clock*, you use the fundamental metaphor *space and movement* to understand that the meeting had an extended duration in time. According to the theory, you create a mental space that (virtually) includes a straight line that represents time. That line is marked with a point representing the beginning of the meeting and another point marking the end of the meeting. Hence, extended duration in time is understood with reference to distance in space, with greater distances corresponding to greater amounts of time. You can also use space and movement to understand financial transactions. If someone says *When the old man died, the inheritance went to John*, you mentally conceptualize the acquisition of wealth by John as an object (*the inheritance*) that moved through space from a starting location (*the old* man) to a goal location (*John*).

We use the fundamental metaphor of *force, agency, and causation* to understand seemingly literal expressions like *Evelina is polite to Ted*, and to understand how the meaning of that expression differs from expressions with very similar meanings, such as *Evelina is civil to Ted* (Pinker, 1994). The difference between *civil* and *polite* in terms of the *force, agency, and causation* metaphor is that the force that causes Evelina to be polite is internal to Evelina, while being *civil* involves a metaphoric situation where there are two forces in opposition, an internal force that is driving Evelina to be mean to Ted and an external force that is preventing her from doing so.

According to the *weak form* of the conceptual-mapping hypothesis, underlying metaphors are not necessary and do not necessarily make up all of a concept's definition, but they are nonetheless routinely called to mind when words are used. So, although the domain *war* does not define the total meaning of the word *argument*, concepts directly related to *war* will become activated when we talk about arguments. Under this formulation, *argument* has its own definition and its own set of relationships between its component concepts, but the way those components are related to one another is influenced by the *war* metaphor. Because the *argument is war* metaphor is commonly used to talk about arguments, the way you mentally define and talk about arguments is shifted to become more similar to the way you mentally define and think about war.

Linguists like George Lakoff also contend that implicit metaphors play a powerful role in political discourse (e.g. Lakoff, 2002, 2008). For instance, when a politician talks about *tax relief* that situates the concept of taxation inside a medical metaphor. Within that metaphorical realm, taxes are equivalent to disease that causes pain. The appropriate response to pain in the medical context is to remove the cause of the pain, and the person who removes that pain is a hero. By saying *tax relief*, a politician can implicitly activate the medical disease metaphor, getting you to think of taxation as disease that causes pain and consider the politician who removes the cause of the pain as the heroic doctor. According to Lakoff, listeners' attitudes can easily be influenced by implicit metaphors, with substantial benefits accruing to politicians who have mastered the use of metaphor.

Although the conceptual-mapping hypothesis has gained popularity, it is not without its critics (Keysar et al., 2000; Murphy, 1996). The strong version of conceptual mapping has been criticized because metaphoric mappings do not appear to be necessary for a concept to be understood, and it may not be possible to work out the correct mappings between different domains without resorting to scientifically unacceptable assumptions. The fact that children may have a well-developed understanding of *anger* well before they develop an understanding of the physics involved in fluid dynamics calls into question the necessity of metaphoric mappings to understand anger.[5]

The conceptual mapping hypothesis also does not explain why some aspects of source domains are included in metaphoric understanding, while others are not. In the *argument is war* mapping, attacking and defending forces, terrain, and methods of attack are used in expressions that utilize the metaphor, but chains of command, logistics, and uniforms, all of which are present in real, literal war situations, never appear in metaphoric comparisons of argument and war. Some of the aspects of the source domain (*war*) will be useful in describing, defining, or thinking about the target concept (*argument*), but many will not. The question, then, is how does the mind know which aspects of the source domain to use to define the target? Unless there is a little person in your mind (a *homunculus*) who already knows how arguments relate to wars, it may be impossible for just the right set of mappings to emerge (see Rai and Chakraverty, 2020, for a computational evaluation of the issues).

Critics of both the strong and weak forms of the conceptual mapping hypothesis also object to the circular nature of some of the evidence mustered in support of the theory. To provide evidence for the conceptual mapping theory, proponents point out commonly used expressions (*I was crushed, She shattered my defenses*, and so on) and then formulate a metaphoric relation that is common across those commonly used expressions (e.g. *argument is war*). Then they provide further examples of commonly used expressions to confirm that people think about arguments by referring to concepts directly related to war. The problem is that the outcome predicted by the theory is the same as the observations that were used to generate the hypothesis in the first place. Ideally, we would like our theories to predict observations not yet obtained. As Keysar and colleagues (Keysar et al., 2000, p. 577) explain, "Using only linguistic evidence for deep connections between language and thought is circular … How do we know that people think of happy and sad in terms of up and down? Because people talk about happy and sad using words such as up and down. Why do people use expressions such as *his spirits rose?* Because people think of happy in terms of UP. Clearly, these arguments are circular and provide no substantive support for the [concept mapping hypothesis]."[6] Further, while there may be close mappings between aspects of the concept argument and aspects of the concept war, and while people may agree that war provides a number of terms that are useful in talking and thinking about arguments, it is always possible that argument and war are represented separately, and that people recognize the

connections between them only after they have been made explicit. Even though the connections between argument and war may seem very compelling, war may still not serve as the fundamental basis of the way we define the word *argument*.

Some critics have also challenged the conceptual mapping hypothesis' assumption that words can have only a single literal meaning. According to conceptual mapping, we would interpret the words *arm*, *leg*, *back*, and *seat* when they refer to parts of a chair by mapping those terms onto the source domain *human body*. There is no logical reason, however, why the same word could not be used to refer literally to a part of a chair and, separately, to a part of the body (just as we reuse words like *bank* to literally refer to the side of a river and a place to keep money).

A final problem with the conceptual-mapping hypothesis is that it does not say what should happen when a target concept is related to several different metaphors. For example, *argument*, besides being metaphorically related to the *war*, is also related to *building* (as in, *that argument needs buttressing, you're on a shaky foundation*), *container* (*she unpacked her claims about metaphor*), and *journey* (*I couldn't follow her line of reasoning, she left me behind when she started talking about idioms*). If a concept already has a definition (*argument is war*), why does it need a second (or third or fourth)? Worse still, some concepts with connections to multiple metaphoric domains inherit contradictory characteristics from different metaphors. As Murphy (1996) explains, *love* is conceptualized as a *journey* in which the participants cooperate to achieve a common goal, and simultaneously as a *valuable commodity* involved in a commercial exchange where the participants are in direct competition in pursuit of different goals.

The structural similarity view

To overcome some of the limitations of the conceptual-mapping view, Murphy (1996) proposed the structural similarity view, which is similar to Dedre Gentner's (1983) *structure-mapping* view. According to these accounts, all concepts are directly represented. For instance, there is a defined concept of "argument" that is represented separately from a separate defined concept of "war." This formulation solves the homunculus problem associated with the conceptual-mapping view. There is no need to propose a little person who makes sure that inappropriate relationships are filtered out when a source domain is used to define another concept. As Murphy (1996, p. 187) proposes, "No one infers that guns are used in arguments, because one already knows that they are not." The structure-mapping view proposes that concepts can be related to one another on the basis of similarity, but that similarity does not make one concept "parasitic" on another. So, we can understand love as a journey, as insanity, as a valuable commodity, and as a number of other things by sequentially relating love to each of these other domains. We can then appreciate similarities between the experience of love and other kinds of experiences, without having self-contradictory elements within the definition of the term *love* itself (as when *journey* and *valuable commodity* suggest both cooperation and competition).[7]

The career-of-metaphor hypothesis

The career-of-metaphor hypothesis represents a hybrid of the comparison and class-inclusion views of metaphor interpretation (Bowdle and Gentner, 2005, 2020). According to this hypothesis, metaphors (metaphorically) have careers, like people have careers.

When you start your career, you normally do the kinds of dirty jobs that nobody else wants to do, but as you progress, your duties and your behaviors change. Likewise, the way metaphors behave and the way they are interpreted change as the metaphor goes from being novel (new) to being frozen (old). According to the career-of-metaphor hypothesis, in the beginning, when a metaphorical expression is first coined, the metaphor is understood by a process of comparison and property matching. But as the metaphor becomes widely used and familiar, comprehenders shift to a category-inclusion mode of processing, as proposed by Glucksberg and colleagues.

Bowdle and Gentner (2005) explain, "as metaphors are conventionalized, there is a shift in mode of processing from comparison to categorization" (p. 194). They note that words commonly used in metaphors, such as *roadblock* or *bottleneck*, started out referring to concrete, real, directly perceivable objects. When people started noticing similarities between some other situation (like a problem at work) and the object referred to literally by words like *roadblock* or *bottleneck*, they would say, *Wow, this problem at work is just like a roadblock*, or *That darned Xerox™ machine is acting like a bottleneck on our project*. At this stage of the expression's career, the critical words *roadblock* and *bottleneck* refer to their literal, directly perceivable real-world concepts. After the metaphoric comparison comes to be commonly used, people drop the *like* and just say, *That problem is a roadblock*. At this stage, the critical words' meaning is based on dual reference (as in the category-inclusion hypothesis), with *roadblock* and *bottleneck* standing for the category *things that stop you from getting where you want to go/things that slow down a process*.

The chief evidence in favor of the career-of-metaphor hypothesis takes the form of aptness ratings for novel versus familiar metaphoric expressions. According to Bowdle and Gentner (2005), novel metaphors are interpreted via a process of comparison, as described in the conversion-to-simile approach. Older, more established metaphors are treated as in the category-inclusion hypothesis, with dual reference and assignment of superordinate category properties of the vehicle to the topic. If new metaphoric expressions are mentally converted to similes, then subjects who are given a choice between a simile (*A is like B*) and a category-inclusion statement (*A is B*) should prefer the simile form for novel comparisons (because it saves them processing step). If established metaphors are treated as category-inclusion statements, then the opposite pattern should hold. For established metaphors, participants should prefer the category-inclusion form over the simile form. Bowdle and Gentner presented subjects with both novel metaphors (e.g. *dancers are butterflies*) and established metaphors (e.g. *problems are roadblocks*) in both the category-inclusion (*A is B*) and simile (*A is like B*) forms. As predicted by the career-of-metaphor hypothesis, people preferred the simile form for novel metaphors, but they preferred the category-inclusion form for established metaphors. Comprehension times for novel metaphoric comparisons were also shorter when they were expressed in simile form than when they were expressed in category-inclusion form.

One potential problem with this line of reasoning is that novel metaphors can be less apt, or less meaningful, than more established metaphors. Novel metaphors survive to become established metaphors, presumably, because they are successful at conveying some useful bit of meaning. Perhaps people prefer the simile form, not because an utterance is novel, but because the utterance is not a particularly apt metaphor. Perhaps people prefer the category-inclusion form over the simile form because the statement in question effectively and efficiently taps a superordinate category, and using the simile form un-aptly focuses attention on the literal, basic category level rather than the more meaningful superordinate category level. To test that hypothesis Sam Glucksberg and Catrinel Haught (Glucksberg and Haught, 2006a, 2006b)

presented people with novel and established metaphors that were equated for aptness ratings. Novel metaphors were rated as equally good and were read and comprehended equally quickly, whether they were expressed in category-inclusion form (e.g. *My lawyer is an old shark*) or simile form (e.g. *My lawyer is like an old shark*). Further, some of the *novel* metaphors were rated as being better when they were expressed in category-inclusion form than when they were expressed in simile form, which poses a problem for the career-of-metaphor hypothesis.

Why Metaphor?

Metaphor is pervasive in everyday life, not just in language but in thought and action. Our ordinary conceptual system, in terms of which we both think and act, is fundamentally metaphorical in nature.
LAKOFF AND JOHNSON

Metaphors are necessary and not just nice.
ANDREW ORTONY[8]

The following conversation between Stephen Colbert (SC), conservative television commentator, and Elizabeth Alexander (EA), Inaugural Poet and Yale University professor, took place on January 21, 2009 (emphases mine):

SC: Let's talk about meaning for a second, OK? Metaphors, OK? *What's the difference between a metaphor and a lie?* Because, you know, "I am the sun. You are the moon." That's a lie, you're not the moon. I'm not the sun, okay? What's the difference between a metaphor and a lie?

EA: Well, that was both a metaphor and a lie. So, the two are not necessarily exclusive. *A metaphor is a way of using language where you make a comparison to let people understand something as it relates to something else*, and that's how we use the language to increase meaning.

SC: *Well, why not just say what you mean, instead of dressing things up in all this flowery language* like, you know, the great romantic poets, you know "Shall I compare thee to a summer's day?" Why not just say "You are hot, let's *do* it"?

There are at least three good answers to Colbert's question.[9] The first has to do with pragmatics and the social nature of speech. The second and third have to do with communicative efficiency.

First, although we can just come right out and say what we are thinking in a very direct way, that can be risky in many social situations. You don't always want your boss to know what you think. So you can either lie, try to escape the conversation, or produce an utterance with some kind of double meaning, like a metaphor.[10] The risk to one's ego and well-being grows exponentially when the topic is love. The problem with using very direct language in affairs of the heart is the crushing loss of face that can result when a direct approach is rejected. So, instead of saying, *You are hot, let's do it* you might say, *If you're not doing anything Saturday night, I know a great little restaurant*. That way, if the indirect approach fails, there is no loss of face, because there was no direct request for a romantic date, and the approacher can act as if nothing important just happened. If you say *You are hot and I would like you to go on a date with me*, and the answer is *no*, there is no easy way to put a positive spin on that.

The second reason why you might choose a metaphor over a more literal form is because metaphors can pack a huge amount of complex meaning into a very small, tight package. The Barenaked Ladies could have said, *It would not bother me if you accompanied me everywhere all the time, even if you irritated all of my friends and co-workers, made it much more difficult for me to do my job, and therefore caused my friends and co-workers to abandon me, reduced my artistic output, and lowered my earnings.* Instead, they just say *You can be my Yoko Ono.* Assuming that you have the background knowledge about Yoko Ono's relationship with John Lennon, which played a major role in breaking up the Beatles, you can easily map the novel situation onto the familiar situation without having to have the entire list of inferences made explicit for you. (Pop quiz: If you know someone who is your Yoko Ono, who are you, according to the metaphor?[11])

Third, good, apt metaphors help the listener make sense of the new information that the speaker wishes to convey. The resulting interpretation is therefore likely to be more accurate and more memorable. The benefit of apt metaphors is a result of mappings between a new and unfamiliar domain to an older, more familiar, and better understood domain. Comprehenders can use the well-understood domain to organize their understanding of the new information, and the information associated with the known domain can also serve as retrieval cues for the new information. This is a tried-and-true strategy for teachers and fiction writers, as exemplified by the *Futurama* episode "Where No Fan Has Gone Before":

Fry: Well, usually on [*Star Trek*] someone would come up with a complicated plan then explain it with a simple analogy.

For instance, when they try to defeat the bad guy, Mellvar, the plan is first described without a helpful metaphor:

Leela: If we can re-route engine power through the primary weapons and reconfigure them to Melllvar's frequency, that should overload his electro-quantum structure.

Because the average person has zero chance of understanding that, the scriptwriter provides a helpful metaphor:

Bender Bending Rodriguez: Like putting too much air in a balloon!

Which is easily understood.

Fry: Of course! It's so simple!

Laboratory research confirms the value of good metaphors. Subjects who read texts that use metaphoric expressions comprehended and remembered those texts better than texts where more literal expressions were used to convey the same information (Albritton et al., 1995). In one such study, participants read an expository text that described attempts to reduce crime. Two groups of participants were tested. One of the groups read a text that was based on the metaphor *crime is disease.* The other group of participants read a more literal version of the same text. So, in the *crime is disease* version, the text said *The sources of crime were diagnosed. Officials desperately sought a cure.* The more literal version said *The sources of crime were studied. Officials desperately sought a solution.* After the participants were done reading the text, they were asked to complete a surprise memory test. Subjects were asked to recall individual sentences from the text, and the preceding sentence was used as a recall cue. If you were a subject in the study, you would be given the recall cue *The sources of crime were diagnosed* (from the metaphor version), or

the cue *The sources of crime were studied* (from the literal version), and your task would be to recall the sentence that followed. Participants in the study performed better on the memory test when they had read the *crime is disease* metaphor version than when they had read the more literal version. Being able to map the novel domain (the causes of crime) onto a more familiar domain (disease) allowed participants to build a more coherent and more tightly interconnected mental representation of the expository text, which led to better comprehension and recall. So, when you want someone to understand something better, it helps to provide them with an apt metaphor.

Metonymy and Under-Specification

Metaphors of the form *A* is *B* are fairly obvious and apparent uses of nonliteral language, but there are nonliteral forms that are not quite so obvious that nonetheless turn up frequently and have to be dealt with in our everyday experience of language comprehension. One such form is *metonymy. Metonymy* or *metonymic expressions* occur when a word that normally refers to one thing is used to refer to something else that bears a relationship to that word. An example will help illustrate. When someone says *I spent the weekend reading Dickens, Dickens* is used as a metonym, and the expression as a whole is an example of *metonymy*. Contrast *I read Dickens* to *In 1870, my great-great granny on my mother's side met Dickens*. In this latter case, *Dickens* is being used to refer to the literal human being Charles Dickens. So, names like *Dickens* can be interpreted in at least two ways. They can be interpreted as referring to real live, literal people, or they can refer to products created by those real live, literal people. This creates a challenge for the language-interpretation system. If Dickens were interpreted literally in the expression *I read Dickens*, that should result in people getting very confused. You can't, after all, literally read a human being the way you literally read a book, by turning pages and moving your eyes, and so forth. But people don't get confused. Why is that? And how does the language-processing system know which interpretation to apply when either the literal or metonymic meanings of Dickens could be appropriate at any given point?

Steven Frisson and Martin Pickering have come up with a theory that explains how metonymic expressions are interpreted (Frisson and Pickering, 1999, 2001, 2007; McElree, Frisson et al., 2006; see also Frazier and Rayner, 1990). Their theory deals with *producer-for-product* metonymies (*I read Dickens*), but also with other types, such as place-for-event metonymies, as in *The students protested after Vietnam* (compare to more literal *The students protested in Vietnam*), place-for-institution metonymies, as in *I talked to the convent yesterday* (compare to the more literal *I talked at the convent yesterday*), and controller-for-controlled metonymies, as in *Vladimir Putin invaded Ukraine* (compare to the more literal *Vladimir Putin traveled to Ukraine*).

The interpretation of metonymies poses the same kinds of challenges that the interpretation of other metaphors does. The listener has to solve the recognition problem as well as determining the relationship between the literal meaning and the metonymic meaning. As with metaphoric expressions, listeners could apply a *literal meaning first* strategy, whereby they access the literal meaning of the critical element (*Dickens, Vietnam, the convent*) and attempt to integrate that meaning with the context, proceeding to a nonliteral metonymic meaning only when the initial integration process fails. Alternatively, listeners could adopt a *metonymic meaning first* strategy, whereby they bypass the literal meaning in favor of the nonliteral metonymic meaning. Finally, listeners could attempt to compute both the literal and metonymic meanings at the same time, in parallel.

If listeners adopted the literal-first strategy, then metonymic expressions like *My great-great grandmother read Dickens* should be harder to understand than literal expressions like *My great-great grandmother dated Dickens*. If listeners adopted the metonymic-first strategy, then literal expressions should be harder to understand than metonymic expressions. If both meanings are computed in parallel, then the two kinds of expression should be equally easy to understand. Of course, familiar metonymic words, such as *Vietnam*, may come to have two related senses stored in the mental lexicon, one that relates to the physical location, and one that relates to the famous event that took place in that physical location, the Vietnam war. If so, metonymous words like *Vietnam* and *Dickens* may behave like *polysemous words* (words that have two or more unrelated meanings) in that processing a metonymous word may involve the simultaneous activation of multiple meanings. If so, metonymous words like *Vietnam* and *Dickens* should be harder to process than equivalent words that have just one stored meaning.

These predictions were tested in a series of eye-tracking experiments by Frisson and colleagues (Frisson and Pickering, 1999, 2001, 2007; McElree, Pylkkänen et al., 2006). They monitored subjects' eye movements as they read sentences containing different kinds of expressions. Some of the expressions were familiar metonymic words, such as *Vietnam* and *Dickens*. Some of the expressions had unfamiliar metonymic uses, such as *during Finland*. Some of the sentences required readers to access the metonymic meanings, as in *The students protested during Vietnam* (for the familiar metonymy) and *The students protested during Finland*. Other sentences required readers to access the literal meaning, as in *The students visited Vietnam* and *The students visited Finland*. The eye-movement data showed that subjects had a relatively hard time understanding the unfamiliar metonymic expressions (*The students protested during Finland*). Reading times on the critical word *Finland* were longer than the comparable familiar metonymy *Vietnam*. More importantly, familiar metonymic expressions were processed just as quickly as their literal counterparts. So, the metonymic expression *The students protested during Vietnam* was just as easy to process as the literal expression *The students visited Vietnam*. The results therefore indicate that processing difficulty for metonymic expressions is not determined by the fact that the expression is nonliteral rather than literal. Instead, what made expressions in this experiment easy or difficult to process was whether or not the reader could access a familiar meaning.

Frisson and Pickering (1999) suggest that the processing of metonymic expressions is best understood as a form of *semantic under-specification*. *Semantic* under-specification means that, rather than activating a predefined, detailed sense of a word like *Dickens*, when you hear *Dickens* in the context of *I read Dickens*, you initially activate a wide field of concepts associated with *Dickens*, and you subsequently narrow that field to tailor the interpretation to the specific context that *Dickens* appears in (Frisson and Pickering refer to this latter process as the *homing-in stage*). As they explain, "One abstract, underspecified meaning of a word with a familiar metonymic sense and a literal sense is initially activated. This meaning is … the same for both senses. Hence, no extra processing is predicted for either sense" (Frisson and Pickering, 1999, p. 1379). Apart from doing a good job incorporating the reaction-time results, philosophical considerations also favor the under-specification account.

Under-specification helps explain why familiar literal and metonymic expressions are processed equally quickly—the initial under-specified interpretation is compatible with both the literal and nonliteral meanings. The under-specification hypothesis also helps deal with the fact that *Dickens* could refer to any of a large number of associated concepts—all of the books by Dickens, a specific title by Dickens, a specific copy of a specific title by Dickens, a statue of Dickens, a picture of Dickens, and on and on. If we had to activate just a single sense of *Dickens*, it is likely that we would have to engage in a lot of repair processing a lot of the time.

Idioms and Frozen Metaphors

There exists a huge dump of worn out metaphors which have lost all evocative power and are merely used because they save people the trouble of inventing phrases for themselves.
GEORGE ORWELL

Idioms make use of ordinary words that have conventional meanings, but when the words are put together in a phrase, the meaning of the utterance is much greater than the sum of the parts. *Screw the pooch* is an English idiom, and its meaning "to blunder" or "to make a mistake" does not transparently relate to any of the individual words in the expression (e.g. Gibbs et al., 1990; Jackendoff, 1995). Idioms of this type are called *non-decomposable*, because they cannot be broken down into parts smaller than the whole expression. They contrast with *decomposable* idioms, which can be broken down into subparts, each of which can be related to a component of the idiom's definition. *Spill the beans*, as an example, can be broken down into the verb by itself, which maps on to *tell*, and *the beans* which maps on to *secrets*. Because individual words in idioms do not refer to their normal or default meanings, idiom comprehension poses a challenge to the language-comprehension system. Somehow, the system must recognize that words are not being used in their normal sense, and it must recover the meaning that is generally assigned to the idiom as a whole. How does the comprehension system accomplish these tasks?

The classical view of idiom comprehension sees idioms as being long words that are analyzed and interpreted as wholes (Chomsky, 1980; Katz, 1973). These accounts view idioms as being essentially "dead metaphors," which were once analyzed like other novel metaphoric expressions, but which over time have become conventionalized and associated with fixed, stored meanings, similar to other words in your vocabulary. As Gibbs et al. (1989, p. 576) explain, "This view suggests that the figurative meanings of idiomatic phrases are directly stipulated in the mental lexicon much like the meaning of an individual word is listed in a dictionary." According to this view, listeners do *not* access the meanings of individual words within the idiomatic expression as the idiom is being comprehended. Instead, the expression as a whole triggers lexical access for the idiomatic meaning, bypassing the normal syntactic and semantic analysis of the idiom's individual parts. Research on the processing and interpretation of idioms has led to a more nuanced, view, however.

According to the *idiom decomposition* account of idiom comprehension, the way an idiom is processed and interpreted depends on specific details of the idiom (e.g. Gibbs, 1980, 1986; Gibbs and Nayak, 1991; Nayak and Gibbs, 1990; Gibbs et al., 1989). For instance, idioms can differ in terms of their *decomposability*—whether individual words in the idiom are associated with individual aspects of its meanings. Idioms such as *spill the beans* and *pop the question* are viewed as *decomposable*, because individual words can be tied to specific parts of the idiomatic meaning (e.g. *spill* = *tell, beans* = *secret; pop* = ask (suddenly), *question* = *marriage proposal*). The category of decomposable idioms can be further subdivided into *normally* and *abnormally* decomposable subcategories. *Lay down the law* is viewed as normally decomposable because there is a semantic relationship between the conventional meanings of its component words and their idiomatic meanings (law = *rules;* the standard meanings of *law* and *rule* are similar). *Spill the beans* is viewed as abnormally decomposable because the normal definitions of *beans* and *secrets* are semantically unrelated. Non-decomposable idioms such as *screw the pooch* (blunder) and *bury the hatchet* (make peace) cannot be analyzed into subparts that map directly to some subpart of the idiomatic meaning.

Decomposable idioms (e.g. *spill the beans*) and non-decomposable idioms (e.g. *screw the pooch*) behave differently along a couple of dimensions. First, decomposable idioms are more *syntactically flexible* than non-decomposable idioms (Gibbs et al., 1989, 1989). This means that if you rearrange the parts of a decomposable idiom, you are less likely to interfere with the idiomatic meaning than if you rearrange the parts of a non-decomposable idiom. So, if someone says *the question was popped by Ted*, you are likely to view that as having an idiomatic meaning (*propose marriage*). But if someone said, *the pooch was screwed by Ted*, that is far less likely to retain its idiomatic meaning (*blunder*) and is more likely to be assigned a more conventional meaning (e.g. *Ted cheated the pooch*).

Decomposable idioms are also more likely to be *lexically flexible* than non-decomposable idioms. This means that you can replace individual words in the idiom with other words and still retain the idiomatic meaning. So, you can change the decomposable idiom *button your lip* to *fasten your lip, button your gob*, or (for our Scottish friends) *hush a gob* without losing the idiomatic meaning ("be quiet"). There are also differences in the amount of time it takes people to process and interpret decomposable and non-decomposable idioms (Gibbs and Gonzales, 1985; Gibbs et al., 1989).

People can process and understand decomposable idioms faster than non-decomposable idioms. Gibbs suggests that this speed advantage for decomposable idioms means that listeners process words in the idiom individually and assign them individual meanings while building up an interpretation of the idiom as a whole (contrary to the classic "dead metaphor" view; Chomsky, 1980; Katz, 1973). As Gibbs et al. (1989, p. 587) explain, "When an idiom is decomposable subjects can assign independent meanings to its individual parts and will quickly recognize how these meaningful parts combined to form the overall figurative interpretation of the phrase."

This line of research therefore challenges the idea that idiomatic expressions are treated as unanalyzable wholes. Additional research suggests that the conventional meanings of individual words can contribute to the meaning of idioms, even for non-decomposable idioms. For example, Gibbs and colleagues note that, although the origin of non-decomposable idioms may be lost or obscure, there are often reasons why particular words and phrases take on the meanings they convey in contemporary language. For example, Hamblin and Gibbs (1999, p. 35) explain the origin of the phrase *kick the bucket*:

> *[It] originally came from one method of slaughtering hogs where hogs were tied by their feet on a wooden frame, called a "buquet" in French and then had their throats cut with a knife. People commented that, when the hogs died, they "kicked the buquet" … Soon enough, people simply talked about dying for animals and human beings as "kicking the buquet"… using one salient part of a scene to refer to the entire complex situation.*

Why does the expression contain the verb *kick*? Originally because of a literal, directly perceived sequence of events that involved kicking. But later, new situations could be mapped onto the conventional scenario. Critically, the idiom *kicked the bucket* still conveys the idea of a sudden event, because the action of kicking is itself a sudden event.

To show that individual elements of idioms still carry the meaning associated with individual words, Gibbs and colleagues presented subjects with idioms such as *kicked the bucket* (which implies a quick action) and *chewed the fat* (which implies a slower action). These idioms were presented in contexts that were either consistent with the type of action conveyed by the verb (*kicked* or *chewed*) or inconsistent with the type of action conveyed by the verb. So, idioms sometimes appeared in contexts that implied a fast

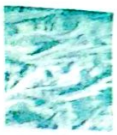

CONSISTENT CONTEXT FOR *CHEW THE FAT*	INCONSISTENT CONTEXT * FOR *CHEW THE FAT*
Joan and Sally are best friends. They have been confiding in each other for years. Every Thursday, they meet for coffee and talk. They often talk for hours while catching up on each other's lives. "You start," said Joan. "Tell me about what has been going on." Later, Joan told her husband about the conversation with Sally. "*We chewed the fat.*"	Joan and Sally are co-workers. Joan overheard some important news as she walked by the boardroom. Their office branch was going to be closing permanently. Joan saw Sally a little while later. "Have you heard the news?" Joan asked. "We might be losing our jobs." Later, Joan told her husband about the conversation with Sally. "*We chewed the fat.*"

(Hamblin and Gibbs, 1999)

action when the idiom itself implied a slow action, and vice versa. People viewed the consistent versions as being more sensible than the inconsistent versions, which suggests that they had greater difficulty integrating the meaning associated with the idiom, in particular the fast or slow action indicated by the verb, when the context and idiom pointed towards different kinds of actions (see also Kessler et al., 2021).

One of the shortcomings of research on idioms is that the vast majority of it has been conducted on English, and there on a restricted range of idiomatic expressions (Kreuz and Graesser, 1991). As a result, conclusions that have been drawn may not reflect universal properties of listeners, and they may reflect idiosyncratic properties of the restricted range of idioms tested in previous studies. Patrizia Tabossi and colleagues tested Italian speakers' responses to a wide range of idiomatic expressions and found evidence that may be incompatible with the idiom decomposition hypothesis, as formulated by Gibbs and colleagues (Tabossi et al., 2008). They argued that the idiom-decomposition hypothesis makes assumptions that may not hold up under closer scrutiny.

First, the decomposition hypothesis assumes that idiomatic expressions convey one specific meaning, and that meaning is captured by one specific literal paraphrase (e.g. *pop the question* equals exactly *propose marriage*, not *ask someone to marry suddenly*, or *kneel and suddenly ask "Will you marry me?"*).[12] Further, it assumes that people generally agree on how the individual parts of idioms connect to individual parts of their meanings. This assumption is based on inter-rater agreement for the kinds of idioms that have been tested in English (about 40 idioms in all).

Tabossi and colleagues noted that languages generally have many more than 40 idiomatic expressions, and they set about testing whether the idiom decomposition assumptions held for much larger samples of idioms. To do so, Tabossi had naive Italian raters assess Italian idioms (e.g. *tirare la caretta/pull the two-wheel cart*, "live a difficult life"; *essere al fresco/be at the fresh*, "be in jail"; *prendere un granchio/take a crab*; "make a blunder"). More specifically, subjects judged whether the idioms were decomposable, and, if so, whether they were normally or abnormally decomposable (as in the earlier studies by Gibbs and colleagues). The results showed that subjects did *not* agree on the choice of

whether an idiom should be considered decomposable or, for those idioms that got the highest overall decomposable ratings, which ones were normally versus abnormally decomposable.

Recall also that the idiom decomposition hypothesis predicts that decomposable idioms should be more syntactically flexible than non-decomposable idioms. When Tabossi and colleagues tested their different kinds of idioms, they found that each type was about equally affected by different kinds of syntactic changes. They also failed to find a speed advantage for any kind of idiom relative to the others (see also Titone and Connine, 1999), although they did replicate previous findings showing that the nonliteral meanings of familiar idioms became available faster than their literal meanings (as expressed in literal paraphrases of the idioms). They concluded therefore that only a very restricted range of idioms behave in regular ways, as the idiom-decomposition hypothesis predicts.

Results like these have motivated researchers like Tabossi and Christina Cacciari to propose the *configuration hypothesis* to explain how idioms are represented in long-term memory and comprehended online (Cacciari et al., 2007; Cacciari and Tabossi, 1988; Tabossi et al., 2005, 2008). In their view, "idioms are mentally represented as configurations of lexical items without any separate representation in the lexicon" (Cacciari et al., 2007, p. 419). So, there is no prestored phrase-length item that corresponds to the idiom. Instead, the words that make up an idiom are processed in the normal way until the listener has received enough information that the idiom can be recognized as being a familiar configuration of words, or until it is highly likely that the string of words will be completed as a familiar idiom (so, listeners can anticipate the presence of an idiom before the entire idiom is heard; Titone and Connine, 1994). According to this view, idioms have a recognition point just like words do. As a result, idioms that are highly predictable in context, or that have earlier recognition points (perhaps because they are less syntactically flexible or non-decomposable), are processed more quickly than less predictable idioms; and idioms that have earlier recognition points are processed and understood more quickly than idioms that have later recognition points.

Embodiment and the Interpretation of Nonliteral Language

Gibbs also argues that embodiment and mental simulation play an important role in metaphor interpretation (Gibbs, 2003; Gibbs and Colston, 1995; Lakoff and Johnson, 1980a). This conclusion is based on metaphoric expressions that call upon our experience with physical bodies and physical processes that operate on those bodies. For example, the way we talk about emotion frequently evokes the metaphor *emotion is liquid in a container*. In studies where people are asked to talk about their emotional experiences, they oftentimes use terms that refer literally to liquids and physical processes.

One student who was asked to talk about *anger*, spontaneously produced the following (Gibbs, 2003, p. 5, my emphases): "At first, anger *burns* in my chest … the anger just *boiled* inside me. I wanted to grab my boyfriend by his shirt, pin him up against the wall and yell at him for being so stupid … Simply telling him that I was upset made my anger *fizzle out* a little. As we talked, my anger *melted* away." Conventional expressions such as, *blow your stack, flip your lid, get hot under the collar, blow off steam*, and *explode* are also commonly used to express the concept and experience of *anger*. This underlying physical metaphor can color our understanding and experience of emotions like anger.

Cognitive linguists, like Ray Gibbs, refer to the physical processes associated with liquids, heat, and containers, as the *source domain*. They suggest that understanding complex domains like emotion involves mappings between the source domain (heated fluids) and the target domain (emotion).

When you adopt the *anger is heated fluid in a closed container* metaphor, a number of conclusions logically follow. First, the fluid will increase in pressure as more heat is applied. Second, when the container fails, it will do so suddenly. Third, the container will fail without anyone consciously willing or wanting it to fail, so the response is essentially beyond anyone's control. When Gibbs (1992, 2003) assessed people's understanding of anger and their understanding of the physics of heated fluids and containers, he found them to be characterized in similar ways. Further, when people read idiomatic expressions, such as *John blew his stack*, they responded faster to subsequent target words, such as *heat*, which expressed part of the physical characteristics of the source domain. So, when we hear idiomatic expressions, they appear to automatically activate knowledge of the source domain, and we understand that when someone *blows their top* or *flips their lid* they did not mean to do so (just as a pressurized container does not willfully explode, it just explodes).[13] As with word meanings, idiomatic expressions and visuomotor representations also appear to be connected (Wilson and Gibbs, 2007). When participants engage in body actions like swallowing, grasping, and chewing, they respond faster to idiomatic expressions like *he swallowed his pride, she grasped the truth*, and *they chewed on the idea*. If understanding idioms involved only the manipulation of abstract mental symbols there is no particular reason why engaging in body actions should facilitate their interpretation.

The Neural Basis of Nonliteral Language Interpretation

Neurophysiological (ERP) and neuroimaging (fMRI, PET) experiments show that different networks of brain areas respond differently to literal and nonliteral language (e.g. Bottini et al., 1994; Canal and Bambini, 2021; Canal et al., 2017; Citron et al., 2019; Coulson and van Petten, 2002; Eviatar and Just, 2006; Mariana et al., 2021). These studies supplement data from studies of patients with lesions (brain damage), who may have selective deficits in understanding or producing either literal or nonliteral language (e.g. Brownell, 1984; Brownell and Stringfellow, 1999; Winner and Gardner, 1977).

Early research on brain-damaged patients suggested that the right hemisphere of the brain might play a special role in the interpretation of nonliteral speech. Patients with right-hemisphere damage were less able than patients with left-hemisphere damage in matching verbal metaphors (e.g. *he had a heavy heart*) with pictures that described either the nonliteral meaning (e.g. a picture of a sad man) or pictures that described the literal meaning (e.g. a man carrying around a large, heavy heart). Subsequent studies (e.g. Brownell et al., 1990) suggested that right-hemisphere patients' problem was really with the picture-matching task, however, rather than metaphoric language per se (people with right-hemisphere brain damage often exhibit problems dealing with visuospatial information, as indicated by syndromes like *neglect*, where people have problems perceiving and imagining the left half of the visual world). When the task was changed from picture matching to giving a verbal definition of nonliteral expressions, patients with right-hemisphere damage performed much better, indicating that they did know the meanings of many nonliteral expressions. So, right-hemisphere damage does not inevitably produce profound impairment of figurative language understanding (see Thoma and Daum, 2006).

The patient data do not provide a firm basis for making inferences about right-hemisphere contributions to nonliteral language processing. As a result, researchers have turned to neurophysiological and neuroimaging studies of healthy individuals to find out how the brain responds to nonliteral language. Let's consider a couple of the prominent theories that relate brain function to the interpretation of nonliteral language.

One major idea, the *right hemisphere hypothesis*, proposes that, while the left hemisphere dominates the process of analyzing and interpreting literal language, the right hemisphere dominates the process of analyzing and interpreting nonliteral language. This hypothesis received support from the early work on brain-damaged patients as well as the very first neuroimaging research on metaphor comprehension (Bottini et al., 1994). In Bottini and colleagues' study, six individuals had their cerebral blood flow measured as they read sentences that expressed novel metaphoric meanings (e.g. *The investors were squirrels collecting nuts*) or literal meanings (e.g. *The boy used stones as paperweights*). After they read each sentence, participants judged whether it made sense on its literal reading (in which case, the response to the metaphoric sentences should be *no*, while the response to the literal sentences should be *yes*). When blood flow changes in response to metaphoric sentences were compared to blood flow changes in response to literal sentences, a number of right-hemisphere regions showed greater response to the metaphoric sentences (compared to the literal), but no left-hemisphere regions showed similar greater response to metaphoric sentences (see Plate 14). These data were interpreted as showing that the right hemisphere was especially activated by metaphoric language, consistent with the right hemisphere hypothesis.

Subsequent imaging studies have not supported the right hemisphere hypothesis, however. A couple of fMRI experiments by Alexander Rapp and colleagues (Rapp et al., 2004, 2007) involved processing of sentences that conveyed relational metaphors in the *An a is a b* form, as in *Die Worte des Liebhabers sind Harfenklange* ("The lover's words are harp sounds"); or literal category-inclusion statements, as in *Die worte des Liebhabers sind Lugen* ("The lover's words are lies"). After reading the sentences, participants judged whether the sentence conveyed a positive or negative message. As shown in Plate 15, when metaphoric sentences were compared to literal sentences, greater activity was observed in the left hemisphere for the metaphoric sentences, but no differences occurred in the right hemisphere.

Additional fMRI experiments involved sentences as stimuli, contrasts of metaphoric with literal materials, and fMRI measurement methods (Eviatar and Just, 2006; Stringaris et al., 2007; Shibata et al. 2007; see also Reyes-Aguilar et al., 2018). They differed with respect to the tasks that subjects performed after reading the stimuli—some experiments involved categorizing the stimuli as metaphoric or literal, some involved a version of the go/no-go paradigm where participants made an overt response only for sentences that were nonsensical. Despite differences in language (Japanese vs. English vs. German), and differences in the secondary task, all of these studies showed greater left-hemisphere activation for metaphoric stimuli compared to literal stimuli, and little or no difference in the right hemisphere for different sentence types. In fact the study on Japanese (Shibata et al., 2007) showed greater right-hemisphere activity for literal sentences than for metaphoric sentences, which runs directly counter to the right hemisphere hypothesis.

This mixed bag of neuroimaging results has motivated some researchers to formulate alternatives to the right hemisphere hypothesis. One such approach is Giora's *graded salience hypothesis* (e.g. Giora, 2003). According to graded salience, differences between the right and left hemisphere are by-products of the kinds of lexical coding that are undertaken by each hemisphere. As suggested by the coarse-coding hypothesis (Beeman,

1998), right-hemisphere lexical representations are more diffuse and have fuzzier boundaries than left-hemisphere lexical representations. As a result, when lexical representations get activated in the right hemisphere, they are more likely to weakly connect to distantly related concepts. This makes right-hemisphere lexical representations well suited to discovering and highlighting the kinds of distant semantic connections that are important in the understanding of novel metaphors. By contrast, the left hemisphere contains more sharply defined lexical representations, and it activates a narrower range of associations in response to individual words. This makes the left hemisphere good at cleanly and sharply activating prestored semantic relationships, which is ideal for the processing of conventional, familiar, well-worn metaphors. The net result, according to graded salience, is that the left hemisphere is good at activating the *salient* meaning of an expression, while the right hemisphere is better at activating *non-salient* meanings. What makes something salient can be a function of frequency (more frequent meanings are more salient), conventionality (more conventional meanings are more salient), or literality (more literal meanings, everything else being equal, are more salient). Critically, the nonliteral meaning of an expression can be more salient than the literal meaning, if the nonliteral meaning is more frequent. This is just what happens with frozen or familiar metaphors such as *iron fist, paper tiger*, and *bad medicine* (now that you've read this chapter). So, the left hemisphere should deal with salient meanings, including familiar metaphors, while the right hemisphere will play a bigger role in the processing of nonsalient meanings, including novel, *un*familiar metaphors.

The graded salience hypothesis receives some support from fMRI experiments on metaphor processing and recent transcranial-magnetic stimulation (TMS) experiments. Two fMRI studies by Mashal and colleagues (Mashal et al., 2005, 2007; see also Ferstl et al., 2008) involved pairs of words that, when taken together, indicate either a literal (e.g. *paper napkin*) or metaphoric (e.g. *paper tiger*) interpretation. Subjects read each pair of words and then made an explicit judgment about them. Specifically, they judged whether the pair were literally related, metaphorically related, or unrelated. All of the stimuli were pretested to see whether they fit into one of four categories: literal, novel metaphors, conventional (familiar) metaphors, or unrelated. The critical comparison was between the novel and conventional metaphor stimuli. According to the right hemisphere hypothesis, both novel and conventional metaphors should fire up the right hemisphere. But according to the graded salience hypothesis, only the novel metaphors should fire up the right hemisphere, because their metaphoric meanings will be less salient than their literal meanings. For the conventional/familiar metaphors, their nonliteral meanings will be more salient than their literal meanings, so the left hemisphere should play the biggest role in their interpretation. As shown in Plate 16, novel metaphors produced greater response than conventional/familiar metaphors in the right hemisphere, as predicted by graded salience. In particular, greater activity was observed in the right posterior superior temporal sulcus and the right inferior frontal gyrus. Novel metaphors also produced greater activity in the left inferior frontal gyrus, an area that was also activated by sentences in the metaphor condition in other fMRI studies.

Pobric et al. (2008) applied a strong magnetic field as people processed word pairs like those used in experiments by Mashal et al. (2005, 2007). TMS temporarily disrupts the activity of populations of neurons directly beneath the point where the TMS is applied. If the right and left hemispheres process different kinds of metaphors in different ways, as suggested by the graded salience hypothesis, then zapping the left and right hemispheres should have different consequences for people who are reading novel and conventional metaphors. So, in this TMS study, subjects read literal word pairs (*paper napkin*), familiar metaphors (*paper tiger*), novel metaphors (*pearl tears*), or unrelated words (*frog napalm*). When subjects had their right hemispheres zapped with TMS, they

had increased difficulty processing the novel metaphors, but the other kinds of expressions were unaffected. Zapping the left hemisphere interfered with conventional/familiar metaphors, but not the other types of expressions. These data are the first known demonstration of a causal relationship between hemispheric function and metaphor-processing ability. Messing with the right or left hemisphere did *not* have a global effect on language processing, nor did it affect metaphoric language globally. Rather, the effects were restricted to particular kinds of metaphoric expressions. Specifically, TMS interfered with the processing of novel metaphors when it was applied over the right hemisphere, but not the left. This result is most compatible with the graded salience hypothesis, and is additional bad news for the right hemisphere hypothesis.

Researchers have also used the *visual hemifield priming paradigm* to study how the two hemispheres respond to nonliteral language. In the visual hemifield priming paradigm, target words are displayed in either the left visual hemifield (to the left of the spot the subject is looking at or *fixating*) or the right visual hemifield (to the right of the spot the subject is looking at). When words are displayed off center in this way, the image of the word is processed either in the left occipital lobe (if the word is displayed to the right of fixation) or in the right occipital lobe (if the word is displayed to the left of fixation). Although the two cerebral hemispheres share information (via the *corpus callosum*, a thick band of fibers that runs horizontally between the two hemispheres), presenting words off center means that the early response to the target word will predominantly reflect activity in the directly stimulated hemisphere. Researchers can manipulate aspects of the context to force either a literal or a nonliteral interpretation of the target word. In the word–word version of the visual hemifield priming paradigm, word pairs induce either a figurative (*stinging insult*) or literal (*stinging bee*) interpretation of the two-word compound. In a sentence version of the task, sentence context determines whether the target word should be assigned a literal or nonliteral meaning. In a study by Anaki et al. (1998) that used the word–word version of the task, only metaphoric meanings were primed in the right hemisphere (people responded faster to *insult* compared to an unrelated control word after the prime word *stinging; bee* was not primed in either the literal or control conditions). As with the neuroimaging results, however, the visual hemifield priming paradigm does not always show a right-hemisphere advantage for nonliteral meanings. In follow-on studies using the word–word version (Kacinik and Chiarello, 2003) and the sentence version of visual hemifield priming (Faust and Weisper, 2000; Kacinik and Chiarello, 2007), both hemispheres showed priming for both literal and nonliteral meanings, and sometimes the left hemisphere showed bigger priming effects than the right hemisphere for nonliteral meanings (contra the right hemisphere hypothesis).

ERP (evoked response potential) results have been similarly mixed. Using literal word pairs (e.g. *ripe fruit*) and metaphoric word pairs (e.g. *conscience storm*), Arzouan, Goldstein, and Faust (2007) showed that the N400 component was larger for novel metaphors, next largest for familiar metaphors (e.g. *iron fist*), and smallest for literal expressions. They also showed that the N400 was largest over the right side of the brain, which they interpreted as indicating that the brain regions that gave rise to the N400 effects were in the right half of the brain (but, as indicated previously, these kinds of assumptions are shaky when we are dealing with ERP data). Another recent ERP study using the sentence version of visual hemifield priming produced a different outcome. It showed that the brain's electrical response to literal and metaphoric meanings was about the same in both hemispheres—although the results did confirm that metaphoric sentences produced a greater N400 effect than literal sentences in both hemispheres (see also Blasko and Kazmerski, 2006; Kazmerski et al., 2003; Tartter et al., 2002).

Summary and Conclusions

In this chapter, we have looked at strengths and (mostly) weaknesses of the *standard pragmatic view* of nonliteral language processing. Although different neural networks appear to be involved to different degrees in the comprehension and production of literal and nonliteral language, interpretation of nonliteral language does not wait for the failure of literal interpretation. Instead, nonliteral meanings appear to be directly accessible for large classes of nonliteral expressions. Considerable theorizing and research has gone into understanding how metaphors are understood. *Comparison* views of metaphor interpretation, and especially the *salience imbalance* version, do not account very well for a number of aspects of metaphor interpretation. As a result, many language scientists prefer a version of the *class-inclusion hypothesis* (Glucksberg, 1998; Glucksberg and Keysar, 1990; Glucksberg and McGlone, 1999). Class inclusion helps to unify the processing of both literal and metaphoric language by proposing that relational metaphors and literal class-inclusion statements are interpreted in the same way. Specifically, the topic is asserted to be a member of the category exemplified by the vehicle. Other theorists, including Gibbs (e.g. Gibbs, 1994), have attempted to close the gap between literal and nonliteral language processing in the area of idiom comprehension. According to Gibbs, components of idioms retain some of their standard meanings, and this shows up in the way idiomatic meanings fit into their contexts. Other accounts, such as Cacciari and Tabossi's, draw parallels between the processing of regular words and idioms by proposing that both involve multiple meaning activation prior to a recognition point, after which a single stored meaning is rapidly accessed and assigned to the idiomatic expression. Research on the neural basis of language is still in its infancy and is, as a result, what Berkeley Breathed might call "higgledy piggledy." But there is certainly hope there, however, as researchers have some well-articulated theories to work from and what Jeff Spicoli would surely call an "ultimate set of tools."

TEST YOURSELF

1. Try to talk to someone for five minutes without using nonliteral language. Count how many times you slip up and use a metaphor.
2. Take the following idioms, classify them as decomposable or non-decomposable. Rearrange the order of the words in some way. Ask a friend to rate whether the idiomatic meaning is preserved. See if you can predict differences between different kinds of idioms.

   barking up the wrong tree
   a chip on your shoulder
   a piece of cake
   pulling my leg
   shaving his cow
   give him the slip

Notes

1 *Before* is itself a spatial metaphor that we use to talk about time. If you stand *before* a judge, that means you are in front of the judge (as opposed to being behind him or her).
2 In fact, the nonliteral interpretations of some familiar idiomatic expressions may be recovered more quickly than their literal meanings (Gibbs, 1980, 1986).
3 Especially if you are Tom Cruise, although he appears to be doing well these days.
4 Spoiler: If Vladimir Putin is Hitler, and Joe Biden is Saddam Hussein's arch rival, then Biden is either Franklin Delano Roosevelt or Winston Churchill, because FDR and Churchill were Hitler's arch rivals. Puyin:Hitler::Biden:FDR/Churchill.
5 Gibbs might suggest that the understanding of anger as fluid in a container is based on direct experience of our own internal bodily fluids (as in Gibbs, 2001; Gibbs et al., 2004), rather than perceptual experience with other kinds of fluids and containers outside the body. He claims (Gibbs, 2001, p. 6), "Under stress, people experience the feeling of their bodily fluids becoming heated." Personally, I have never experienced the sensations that he attributes to fluids in the body in association with the experience of anger, and even if I did, there is no way for me to tell whether you have had the same subjective experience as me. Further, if we did have those subjective experiences, they could easily be influenced by culturally transmitted expectations, by the very expressions (*he flipped his lid, she boiled over*) that the subjective experiences are meant to give rise to.
6 Recent attempts have been made to muster evidence for the conceptual mapping hypothesis that do not depend on collecting examples of utterances that instantiate the metaphor in question (e.g. Teuscher et al., 2008). The author's view of this evidence is that it offers weak support at best for the strong form of the conceptual mapping hypothesis. Stay tuned.
7 Gibbs responds that complex metaphoric relations, such as that that links *argument* and *war*, can be built up by combining more basic, primitive metaphoric domains (Gibbs et al., 2004). For example, the *arguments are buildings* metaphor consists of a combination of two more primitive metaphors, *persisting is staying upright* and *structure is physical structure*. However, as with the conceptual mapping hypothesis more generally, it is incumbent on the proponents of such theories to explain how just the right combinations of primitive metaphors can emerge without a homunculus.
8 Quotation brazenly stolen from Bowdle and Gentner (2005, p. 193).
9 Ignore for the moment that *You are hot* is an attributional metaphor.
10 As in Monty Python's "splunge" sketch (Python, 1998).
11 You are John Lennon, and you were fabulous on "Live in New York City."
12 Gibbs and colleagues also tend to assign very minimalist standard interpretations to idiomatic expressions. For example, *flip your lid* is assigned the meaning *become angry*, which probably does not match the meaning that people actually assign to that idiom, namely, *become extremely angry very suddenly and uncontrollably* (Kreuz and Graesser, 1991). If people are assumed to have these more complex understandings of idiomatic expressions, then much of the evidence in favor of the idiom decomposition hypothesis evaporates. For example, the fact that people do not like the expression *flipped his lid* in the context of a scenario where someone is slowly becoming angry could just mean that the conventional meaning of the idiom ("suddenly become angry") does not fit, the same way the literal expression *he smiled a big smile and laughed* would not fit.
13 One of the problems that anger-management counselors have to overcome is people's idea that, like pressurized containers, they cannot control and therefore do not bear responsibility for their outbursts.

References

Albritton, D., McKoon, G., and Gerrig, R. (1995). Metaphor-based schemas and text comprehension: Making connections through conceptual metaphor. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*, 612–625.

Anaki, D., Faust, M., and Kravetz, S. (1998). Cerebral hemispheric asymmetries in processing lexical metaphors. *Neuropsychologia*, *36*, 691–700.

Arzouan, Y., Goldstein, A., and Faust, M. (2007). Brainwaves are stethoscopes: ERP correlates of novel metaphor comprehension. *Brain Research*, *1160*, 69–81.

Bambini, V., Bott, L., and Schumacher, P.B. (2021). It is not always a matter of time: Addressing the costs of metaphor and metonymy through a speed-accuracy trade-off study. *Canadian Journal of Experimental Psychology/Revue Canadienne de Psychologie Expérimentale*, *75*(2), 189.

Beeman, M. (1998). Coarse semantic coding and discourse comprehension. In M. Beeman and C. Chiarello (Eds.), *Right Hemisphere Language Comprehension: Perspectives from Cognitive Neuroscience* (pp. 255–284). Erlbaum.

Blasko, D.G., and Connine, C.M. (1993). Effects of familiarity and aptness on metaphor processing. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 19*, 295–308.

Blasko, D.G., and Kazmerski, V.A. (2006). ERP correlates of individual differences in the comprehension of nonliteral language. *Metaphor and Symbol, 21*, 267–284.

Bottini, G., Corcoran, R., Sterzi, R., et al. (1994). The role of right hemisphere in the interpretation of figurative aspects of language: A positron emission tomography study. *Brain, 117*, 1241–1253.

Bowdle, B.F., and Gentner, D. (2005). The career of metaphor. *Psychological Review, 112*, 193–216.

Bowdle, B.F., and Gentner, D. (2020, December). Metaphor comprehension: From comparison to categorization. In *Proceedings of the Twenty First Annual Conference of the Cognitive Science Society* (pp. 90–95). Psychology Press.

Brownell, H. (1984). Sensitivity to lexical denotation and connotation in brain damaged patients. *Brain and Language, 22*, 253–264.

Brownell, H., and Stringfellow, A. (1999). Making requests: Illustrations of how right-hemisphere brain damage can affect discourse production. *Brain and Language, 68*, 442–465.

Brownell, H.H., Simpson, T.L., Bihrle, A.M., et al. Appreciation of metaphoric alternative word meanings by left and right brain-damaged patients. *Neuropsychologia, 28*, 375–383.

Cacciari, C., Padovani, R., and Corradini, P. (2007). Exploring the relationship between individuals' speed of processing and their comprehension of spoken idioms. *European Journal of Cognitive Psychology, 19*, 417–445.

Cacciari, C., and Tabossi, P. (1988). The comprehension of idioms. *Journal of Memory and Language, 27*, 668–683.

Canal, P., and Bambini, V. (2021). *Pragmatics electrified,* doi: 10.31234/osf.io/5z7b2.

Canal, P., Pesciarelli, F., Vespignani, F., et al. (2017). Basic composition and enriched integration in idiom processing: An EEG study. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 43*(6), 928.

Chomsky, N. (1980). *Rules and Representations.* Columbia University Press.

Citron, F.M., Cacciari, C., Funcke, J.M., et al. (2019). Idiomatic expressions evoke stronger emotional responses in the brain than literal sentences. *Neuropsychologia, 131*, 233–248.

Clark, H.H. (1996). *Using Language.* Cambridge University Press.

Clark, H.H., and Lucy, P. (1975). Understanding what is meant from what is said: A study in conversationally conveyed requests. *Journal of Verbal Learning and Verbal Behavior, 14*, 56–72.

Colston, H.L., and Gibbs, R.W. Jr. (2002). Are irony and metaphor understood differently? *Metaphor and Symbol, 17*, 57–80.

Coulson, S., and van Petten, C. (2002). A special role for the right hemisphere in metaphor comprehension? ERP evidence from hemifield presentation. *Brain Research, 1146*, 128–145.

Eviatar, Z., and Just, M.A. (2006). Brain correlates of discourse processing: An fMRI investigation of irony and conventional metaphor comprehension. *Neuropsychologia, 44*, 2348–2359.

Faust, M., and Weisper, S. (2000). Understanding metaphoric sentences in the two cerebral hemispheres. *Brain and Cognition, 43*, 186–191.

Ferstl, E.C., Neumann, J., Bogler, C., et al. (2008). The extended language network: A meta-analysis of neuroimaging studies on text comprehension. *Human Brain Mapping, 29*, 581–593.

Fogelin, R.J. (1988). *Figuratively Speaking.* Yale University Press.

Frazier, L., and Rayner, K. (1990). Taking on semantic commitments: Processing multiple meanings vs. multiple senses. *Journal of Memory and Language, 29*, 181–200.

Frisson, S., and Pickering, M.J. (1999). The processing of metonymy: Evidence from eye movements. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 25*, 1366–1383.

Frisson, S., and Pickering, M.J. (2001). Obtaining a figurative interpretation of a word: Support for underspecification. *Metaphor and Symbol, 16*, 149–171.

Frisson, S., and Pickering, M.J. (2007). The processing of familiar and novel senses of a word: Why reading Dickens is easy but reading Needham can be hard. *Language and Cognitive Processes, 22*, 595–613.

Gentner, D. (1983). Structure mapping: A theoretical framework for analogy. *Cognitive Science, 7*, 155–170.

Gibbs, R.W. Jr. (1980). Spilling the beans on understanding and memory for idioms. *Memory and Cognition, 8*, 449–456.

Gibbs, R.W. (1983). Do people always process the literal meanings of indirect requests? *Journal of Experimental Psychology: Learning, Memory, and Cognition, 9*, 524–533.

Gibbs, R.W. Jr. (1986). Skating on thin ice: Literal meaning and understanding idioms in conversation. *Memory and Cognition, 14*, 149–156.

Gibbs, R.W. Jr. (1992). Why do idioms mean what they do? *Journal of Memory and Language, 31*, 485–506.

Gibbs, R.W. Jr. (1994). *The Poetics of Mind*. Cambridge University Press.

Gibbs, R.W. Jr. (2001). Evaluating contemporary models of figurative language understanding. *Metaphor and Symbol, 16*, 317–333.

Gibbs, R.W. Jr. (2003). Embodied experience and linguistic meaning. *Brain and Language, 84*, 1–15.

Gibbs, R.W. Jr., and Colston, H.L. (1995). The cognitive psychological reality of image schemas and their transformations. *Cognitive Linguistics, 6*, 347–378.

Gibbs, R.W. Jr., and Gonzales, G. (1985). Syntactic frozenness in processing and remembering idioms. *Cognition, 20*, 243–259.

Gibbs, R.W. Jr., Lima, P L.C., and Francozo, E. (2004). Metaphor is grounded in embodied experience. *Journal of Pragmatics, 36*, 1189–1210.

Gibbs, R.W. Jr., and Nayak, N.P. (1991). Why idioms mean what they do. *Journal of Experimental Psychology: General, 120*, 93–95.

Gibbs, R.W., Nayak, N.P., and Cutting, C. (1989). How to kick the bucket and not decompose: Analyzability and idiom processing. *Journal of Memory and Language, 28*, 576–593.

Gibbs, R.W. Jr., Nayak, N.P., and Cutting, C. (1990). Taking on semantic commitments: Processing multiple meanings vs. multiple senses. *Journal of Memory and Language, 29*, 181–200.

Giora, R. (2003). *On Our Mind: Salience, Context, and Figurative Language*. Oxford University Press.

Giora, R. (2007). Is metaphor special? *Brain and Language, 100*, 111–114.

Glucksberg, S. (1998). Understanding metaphors. *Current Directions in Psychological Science, 7*, 39–43.

Glucksberg, S. (2003). The psycholinguistics of metaphor. *Trends in Cognitive Sciences, 7*, 92–96.

Glucksberg, S., Gildea, P., and Bookin, H. (1982). On understanding nonliteral speech: Can people ignore metaphors? *Journal of Verbal Learning and Verbal Behavior, 21*, 85–98.

Glucksberg, S., and Haught, C. (2006a). On the relation between metaphor and simile: When comparison fails. *Mind and Language, 21*, 360–378.

Glucksberg, S., and Haught, C. (2006b). Can Florida become *like* the next Florida? When metaphoric comparisons fail. *Psychological Science, 17*, 935–938.

Glucksberg, S., and Keysar, B. (1990). Understanding metaphorical comparisons: Beyond similarity. *Psychological Review, 97*, 3–18.

Glucksberg, S., Keysar, B., and McGlone, M.S. (1992). Metaphor understanding and accessing conceptual schema. *Psychological Review, 99*, 578–581.

Glucksberg, S., Manfredi, D.A., and McGlone, M.S. (1997). Metaphor comprehension: How metaphors create new categories. In T.B. Ward, S.M. Smith, and J. Vaid (Eds.), *Creative Thought: An Investigation of Conceptual Structures and Processes* (pp. 327–350). American Psychological Association.

Glucksberg, S., and McGlone, M.S. (1999). When love is not a journey: What metaphors mean. *Journal of Pragmatics, 31*, 1541–1558.

Glucksberg, S., McGlone, M.S., and Manfredi, D. (1997). Property attribution in metaphor comprehension. *Journal of Memory and Language, 36*, 50–67.

Grice, H.P. (1989). *Studies in the Way of Words*. Harvard University Press.

Hamblin, J.L., and Gibbs, R.W. Jr. (1999). Why you can't kick the bucket as you slowly die: Verbs in idiom comprehension. *Journal of Psycholinguistic Research, 28*, 25–39.

Harris, R. (1976). Comprehension of metaphors: A test of the two-stage processing model. *Bulletin of the Psychonomics Society, 8*, 312–314.

Holyoak, K.J., and Stamenković, D. (2018). Metaphor comprehension: A critical review of theories and evidence. *Psychological Bulletin, 144*(6), 641.

Inhoff, A.W. (1984). Two stages of word processing during eye fixations in the reading of prose. *Journal of Verbal Learning and Verbal Behavior, 23*, 612–624.

Jackendoff, R. (1995). The boundaries of the lexicon. In M. Everaert, E J. van der Linden, A. Schenk, et al. (Eds.), *Idioms: Structural and Psychological Perspectives* (pp. 133–165). Erlbaum.

Johnson, M.G., and Malgady, R.G. (1979). Some cognitive aspects of figurative language: Association and metaphor. *Journal of Psycholinguistic Research, 8*, 249–265.

Kacinik, N.A., and Chiarello, C. (2003). An investigation of hemisphere differences for moderately imageable words across high and low image contexts. *Brain and Cognition, 53*, 239–242.

Kacinik, N.A., and Chiarello, C. (2007). Understanding metaphors: Is the right hemisphere uniquely involved? *Brain and Language, 100*, 188–207.

Katz, J. (1973). Compositionality, idiomaticity, and lexical substitution. In S. Anderson and P. Kiparsky (Eds.), *A Festschrift for Morris Halle* (pp. 357–376). Holt, Rinehart, and Winston.

Kazmerski, V., Blasko, D., and Dessalegn, B. (2003). ERP and behavioral evidence of individual differences in metaphor comprehension. *Memory and Cognition, 31*, 673–689.

Kessler, R., Weber, A., and Friedrich, C.K. (2021). Activation of literal word meanings in idioms: Evidence from eye-tracking and ERP experiments. *Language and Speech, 64*(3), 594–624.

Keysar, B. (1989). On the functional equivalence of literal and metaphoric interpretations in discourse. *Journal of Memory and Language, 28*, 275–285.

Keysar, B., Shen, Y., Glucksberg, S., et al. (2000). Conventional language: How metaphorical is it? *Journal of Memory and Language, 43*, 576–593.

Kreuz, R.J., and Graesser, A.C. (1991). Aspects of idiom comprehension: Comment on Nayak and Gibbs. *Journal of Experimental Psychology: General, 120*, 90–92.

Lakoff, G. (1987). *Women, Fire, and Dangerous Things: What Categories Reveal About the Mind*. University of Chicago Press.

Lakoff, G. (2002). *Moral Politics: How Liberals and Conservatives Think*. University of Chicago Press.

Lakoff, G. (2008). *The Political Mind: Why You Can't Understand 21st-Century American Politics with an 18th-Ccentury Brain*. Viking.

Lakoff, G., and Johnson, M. (1980a). The metaphorical structure of the human conceptual system. *Cognitive Science, 4*, 195–208.

Lakoff, G., and Johnson, M. (1980b). *Metaphors We Live By*. University of Chicago Press.

Malgady, R.G., and Johnson, M.G. (1980). Measurement of figurative language: Semantic feature models of comprehension and appreciation. In R.P. Honeck and R.R. Hoffman (Eds.), *Cognition and Figurative Language* (pp. 239–258). Erlbaum.

Mariana, B., Carolina, L., Valeria, A., et al. (2021). Functional anatomy of idiomatic expressions. *Brain Topography, 34*(4), 489–503.

Martin, A.E., and McElree, B. (2018). Retrieval cues and syntactic ambiguity resolution: Speed-accuracy tradeoff evidence. *Language, Cognition and Neuroscience, 33*(6), 769–783.

Mashal, N., Faust, M., and Hendler, T. (2005). The role of right hemisphere in processing nonsalient metaphorical meanings: Application of principal components analysis to fMRI data. *Neuropsychologia, 43*, 2084–2100.

Mashal, N., Faust, M., Hendler, T., et al. (2007). An fMRI investigation of the neural correlates underlying the processing of novel metaphoric expressions. *Brain and Language, 100*, 115–126.

McElree, B., Frisson, S., and Pickering, M.J. (2006). Deferred interpretations: Why starting Dickens is taxing but reading Dickens isn't. *Cognitive Science, 30*, 181–192.

McElree, B., and Griffith, T. (1995). Syntactic and thematic processing in sentence comprehension: Evidence for a temporal dissociation. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21*, 134–157.

McElree, B., and Griffith, T. (1998). Structural and lexical constraints on filling gaps during sentence comprehension: A time-course analysis. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 24*, 432–460.

McElree, B., and Nordlie, J. (1999). Literal and figurative interpretations are computed in equal time. *Psychonomic Bulletin and Review, 6*, 486–494.

McElree, B., Pylkkänen, L., Pickering, M. J., et al. (2006). A time course analysis of enriched composition. *Psychonomic Bulletin and Review, 13*, 53–59.

Miller, G.A. (1979). Images and models: Similes and metaphors. In A. Ortony (Ed.), *Metaphor and Thought* (pp. 357–400). Cambridge University Press.

Murphy, G.L. (1996). On metaphoric representation. *Cognition, 60*, 99–108.

Nayak, N.P., and Gibbs, R.W. Jr. (1990). Conceptual knowledge in the interpretation of idioms. *Journal of Experimental Psychology, 119*, 315–330.

Ortony, A. (1975). Why metaphors are necessary and not just nice. *Educational Theory, 25*, 45–53.

Ortony, A. (1979). Beyond literal similarity. *Psychological Review, 86*, 161–180.

Pinker, S. (1994). *The Language Instinct*. Harper.

Pobric, G., Mashal, N., Faust, M., et al. (2008). The role of the right cerebral hemisphere in processing novel metaphoric expressions: A transcranial magnetic stimulation study. *Journal of Cognitive Neuroscience, 20*, 170–181.

Pollio, H., Barlow, J., Fine, H., et al. (1977). *Psychology and the Poetics of Growth: Figurative Language in Psychotherapy and Education*. Erlbaum.

Pollio, H.R., Fabrizi, M.S., Sills, A., et al. (1984). Need metaphoric comprehension take longer than literal comprehension? *Journal of Psycholinguistic Research, 13*, 195–214.

Python, M. (1998). *The Complete Monty Python: Vols. 1 and 2*. Methuen.

Rai, S., and Chakraverty, S. (2020). A survey on computational metaphor processing. *ACM Computing Surveys (CSUR), 53*(2), 1–37.

Rapp, A.M., Leube, D.T., Erb, M. et al. (2004). Neural correlates of metaphor processing. *Cognitive Brain Research, 20*, 395–402.

Rapp, A.M., Leube, D.T., Erb, M., et al. (2007). Laterality in metaphor processing: Lack of evidence from functional magnetic resonance imaging for the right hemisphere theory. *Brain and Language, 100*, 142–149.

Reyes-Aguilar, A., Valles-Capetillo, E., and Giordano, M. (2018). A quantitative meta-analysis of neuroimaging studies of pragmatic language comprehension: In search of a universal neural substrate. *Neuroscience, 395*, 60–88.

Searle, J. (1979). Meaning and use. *Studies in Linguistics and Philosophy, 3*, 181–197.

Shibata, M., Abe, J., Terao, A., et al. (2007). Neural bases of metaphor comprehension: An fMRI study. *Cognitive Studies: Bulletin of the Japanese Cognitive Science Society, 14*, 339–354.

Shinjo, M., and Myers, J.L. (1987). The role of context in metaphor comprehension. *Journal of Memory and Language, 26*, 226–241.

Stern, J. (2000). *Metaphor in Context*. MIT Press.

Stringaris, A.K., Medford, N.C., Giampietro, V., et al. (2007). Deriving meaning: Distinct neural mechanisms for metaphoric, literal, and non-meaningful sentences. *Brain and Language, 100*, 150–162.

Stroop, J.R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology, 18*, 643–662.

Swinney, D., and Cutler, A. (1979). The access and processing of idiomatic expressions. *Journal of Verbal Learning and Verbal Behavior, 18*, 523–534.

Tabossi, P., Fanari, R., and Wolf, K. (2005). Spoken idiom recognition: Meaning retrieval and word expectancy. *Journal of Psycholinguistic Research, 34*, 465–495.

Tabossi, P., Fanari, R., and Wolf, K. (2008). Processing idiomatic expressions: Effects of semantic compositionality. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 34*, 313–327.

Tartter, V.C., Gomes, H., Dumbrovsky, B., et al. (2002). Novel metaphors appear anomalous at least momentarily: Evidence from N400. *Brain & Language, 80*, 488–509.

Teuscher, U., McQuire, M., Collins, J., et al. (2008). Congruity effects in time and space: Behavioral and ERP measures. *Cognitive Science, 32*, 563–578.

Thoma, P., and Daum, I. (2006). Neurocognitive mechanisms of figurative language processing: Evidence from clinical dysfunctions. *Neuroscience and Biobehavioral Reviews, 30,* 1182–1205.

Titone, D.A., and Connine, C.M. (1994). Comprehension of idiomatic expressions: Effects of predictability and literality. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20,* 1126–1138.

Titone, D.A., and Connine, C.M. (1999). On the compositional and noncompositional nature of idiomatic expressions. *Journal of Pragmatics, 31,* 1655–1674.

Tourangeau, R., and Sternberg, R.J. (1981). Aptness in metaphor. *Cognitive Psychology, 13,* 27–55.

Tversky, A. (1977). Features of similarity. *Psychological Review, 85,* 327–352.

Wilkes-Gibbs, D., and Clark, H.H. (1992). Coordinating beliefs in conversation. *Journal of Memory and Language, 31,* 183–194.

Wilson, N.L., and Gibbs, R.W. Jr. (2007). Real and imagined body movement primes metaphor comprehension. *Cognitive Science, 31,* 721–731.

Winner, E., and Gardner, H. (1977). The comprehension of metaphor in brain-damaged patients. *Brain, 100,* 719–727.

Wolff, P., and Gentner, D. (2000). Evidence for role-neutral initial processing of metaphors. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 26,* 529–541.

Zhou, Y., and Tse, C.S. (2022). Sweet taste brings happiness, but happiness does not taste sweet: The unidirectionality of taste-emotion metaphoric association. *Journal of Cognitive Psychology, 34*(3), 339–361.