Mental simulation in literal and figurative language understanding[*]

Benjamin Bergen
Department of Linguistics
University of Hawaii at Manoa

1. Introduction

Suppose you ask a colleague how a class he just taught went, and he replies that "It was a great class - the students were glued to their seats." There are two clearly distinct interpretations of this utterance, which might usefully be categorized as a literal and a figurative one. Which interpretation you believe is appropriate will determine your course of action; whether you congratulate your colleague for his fine work, or whether you report him to the dean for engaging in false imprisonment.

What distinguishes between these two types of interpretation? Classically, the notions of literalness and figurativity are viewed as pertaining directly to language - words have literal meanings, and can be used figuratively when specific figures of speech cue appropriate interpretation processes (Katz 1977, Searle 1978, Dascal 1987). Indeed, this approach superficially seems to account for the two interpretations described above. The literal (and one would hope, less likely) interpretation involves simply interpreting each word in the sentence and their combination in terms of its straightforward meaning - the glue is literal glue and the seats are literal seats.

The figurative interpretation, by contrast, requires knowledge of an idiom - a figure of speech. This idiom has a particular form, using the words *glued, to,* and *seat(s)*, with a possessive pronoun or other noun phrase indicating who was glued to their seat(s) in the middle. It also carries with it a particular meaning: the person or people described were in rapt attention. Thus the figurative interpretation of the sentence differs from the literal one in that the meaning to be taken from it is not built up compositionally from the meanings of words included in the utterance. Rather, the idiom imposes a non-compositional interpretation. As a result, the meaning of the whole can be quite distinct from the meanings of the words included within the idiom. This distinction between idiomaticity and compositionality is an important component of the classical figurative-literal distinction. A consequence of equating figurativity with particular figures of speech is that figurative language can be seen as using words in ways that differ from their real, literal meaning.

While this classical notion of figurative and literal language may seem sensible, it leads to a number of incorrect and inconsistent claims about the relation between literal and figurative language. (Note that although it is not strictly correct to talk about language itself as being literal or figurative - instead we should discuss acts of figurative or literal language processing - we will adopt this convention for ease of exposition.) One major problem is that it claims that figurative language processing mechanisms are triggered by language - explicit linguistic cues or figures of speech. Indeed, such appears to be case in the above example, but in many cases it is not. A sentence like *You'll find yourself glued to the latest shareholders' report* can be interpreted "figuratively" even though it has in it no figure of speech equivalent of *glued to X's seat(s)* that might trigger a search for such an interpretation. The only trigger for this figurative meaning is the verb *glue*, which would be hard to justify as an idiom or figure of speech. The absence of overt indicators of figurativity might lead us to hypothesize that some words may encode both literal and figurative meanings (as suggests by

[*] My thanks to Nancy Chang, Seana Coulson, Jerome Feldman, and Michael Israel for useful comments on this work. All errors and omissions are my own.

numerous authors, such as Lakoff and Johnson (1980)). But this polysemy solution is not viable on the original definition of figurativity, where figurative meanings are specifically not encoded in words, but rather are generated by inferential and pragmatic processes. Therefore, if there is a particular meaning associated with the word *glue*, it must be a literal, and not a figurative one. Continuing from the classical perspective, we would thus be forced to conclude that while the most likely interpretation of *My students were glued to their seats* is figurative, the most likely interpretation of *You'll find yourself glued to the latest shareholders' report* is literal - a particularly unparsimonious and dubious claim. The problems with the reasoning outlined here hinge on the notion that figurative language is triggered by linguistic forms and that it involves processes of interpretation beyond the meanings of composite words.

Another problem with the classical view results from the idea that words have particular, literal meanings that are distinct from the figurative uses words can be put to. If linguistic cues, like idioms, are not responsible for triggering figurative language interpretation, then it could be that superficial incoherencies are to blame (Hobbs et al. 1993, Pustejovsky 1991). For example, it might be that figurative interpretation processes used to comprehend sentences like *America was glued to the television as the election scandal unraveled* are triggered by semantic incoherencies: America (as a nation) is not the type of thing that can be literally glued to any physical thing, including a television, nor is an election scandal the type of thing that can physically unravel, not being a fabric.

But this attempt to salvage the classical notion of the literal-figurative distinction is no more successful than the polysemy approach. As argued by various authors (Ortony 1980, Gibbs 1994, Gibbs 1998, Markert & Hahn 1999, Markert & Hahn 2002), figurative language often is ambiguous with semantically coherent literal interpretations. The student-gluing example at the beginning of this paper is a case in point - a literal interpretation (in which the students are actually glued to their seats) is semantically viable, and yet the figurative interpretation is more viable in most contexts.

By the same token, semantic incoherence is not necessarily a trigger for figurative interpretations - sentences like *Colorless green ideas sleep furiously* don't have any particularly accessible figurative interpretation outside of imaginable specific contexts. Similarly, incoherence often triggers processes other than figurative language interpretation, like *construal* (Langacker 1991) also known as *coercion* (Michaelis 2004). In cases of construal, a word is used for a new purpose; a sentence like *When my right rear blew out, the driveway was covered with tire* is incoherent in that *tire* is almost always a count noun, and yet here it is used as a mass noun. In cases of construal like this one, there is no relevant figurative interpretation of the utterance - it is pieces of real, literal, rubber tire that are literally covering the literal driveway. Thus, despite a superficial linguistic incoherence, language using construal is not usefully analyzed as involving figurative language understanding. Examples like these demonstrate once again that associating figurative language with linguistic cues fails to account for the distribution of figurative interpretations of utterances.

There is clearly some substance to the literal figurative distinction, or the current volume would be much shorter than it is. The remainder of this paper proposes a solution to the question of the relation between literal and figurative language, based on a cognitively-oriented alternative to the classical view of literal and figurative meaning. On this theory of linguistic meaning, known as *simulation semantics*, language users construct mental simulations of the perceptual and motor content of experiences described by language they produce or understand, and then propagate inferences on the basis of these mental simulations. As we see below, understanding various types of figurative language can be seen to activate mental simulation of in a variety of ways, with similarities to and differences from the routine use to which this mental function is put during literal language processing.

2. Mental simulation in literal language understanding

Painted in the broadest strokes, approaches to the problem of meaning have followed one of two paths. The first of these sees meaning, in the tradition of objectivist Anglo-American philosophy and functionalist cognitive science (see the discussions in Lakoff 1987 and Barsalou 1999), as the manipulation of abstract and amodal (that is, abstracted away from any modality, such as perception or action) symbols (May 1985). On this view, meanings are static, amodal representations that are associated with words. Understanding language is constituted of determining first which such symbols are denoted by the words in an utterance, and second how they are to be combined. For example, a sentence like *Mary kicked John* would be interpreted by determining which symbolic representations are appropriate for each of the words in the sentence, where *Mary* and *John* are represented by unary symbols, perhaps JOHN and MARY, which contrast with the more complex two-place predicate symbol representing *kicked*, perhaps KICKED(X,Y). These symbols are combined together appropriately, so that the meaning of the sentence is something like KICKED(MARY,JOHN).

An alternative, "embodied" approach views meaning as centrally involving the activation of perceptual, motor, social, and affective knowledge that characterizes the content of utterances. Through exposure to language in context, language users learn to pair chunks of language like *kick, Mary,* or *John* with perceptual, motor, social, and affective experiences. In subsequent instances of language use, when the original perceptual, motor, social, and affective stimuli are not contextually present, the experience of them is re-created through the activation of neural structures responsible for experiencing them in the first place. This view of meaning is "embodied" in that meaning depends on an individual having had experiences in their body in the actual world, where they recreate those experiences in response to linguistic input, and use them to produce meaningful linguistic output.

The notion that language understanding is based on the internal recreation of previous, embodied experiences is supported by recent brain research, showing that motor and pre-motor cortex areas associated with specific body parts (i.e. the hand, leg, and mouth) become active in response to motor language referring to those body parts. Using behavioral and neurophysiological methods, Pulvermüller et al. (2001) and Hauk et al. (2004) found that verbs associated with different effectors were processed at different rates and in different regions of motor cortex. In particular, when subjects perform a lexical decision task (deciding as quickly as possible whether a letter string was a word of their language) with verbs referring to actions involving the mouth (*chew*), leg (*kick*), or hand (*grab*), the motor cortex areas responsible for mouth versus leg versus hand motion received more activation, respectively. Tettamanti et al. (ms.) have also shown through an imaging study that passive listening to sentences describing mouth versus leg versus hand motions activates different parts of pre-motor cortex (as well as other areas).

Behavioral studies also offer convergent evidence for the automatic and unconscious use of perceptual and motor systems during language use. Recent work on spatial language (Richardson et al. 2003, Bergen, Matlock, and Narayanan ms) has found that listening to sentences with visual semantic components can result in selective interference with visual processing. While processing sentences that encode upwards motion, like *The ant climbed*, subjects take longer to perform a visual categorization task in the upper part of their visual field (deciding whether an image is a circle or a square); the same is true of downwards-motion sentences like *The ant fell* and the lower half of the visual field. These results imply that language, like memory, evokes visual imagery that interferes with visual perception.

A second experimental method (Glenberg & Kashak 2002) tests the extent to which motor representations are activated for language understanding. The basic finding of this line of research is that when subjects are asked to perform a physical action, such as moving their hand away from or toward their body, in response to a sentence, it takes them longer to perform the action if it is incompatible with the motor actions described in the sentence. For example, if the sentence is *Andy gave you the pizza*, subjects take longer to push a button requiring them to move their hand away from their body than one requiring them to move their hand towards their body, and the reverse is true for sentences indicating motion away from the subject, like *You gave the pizza to Andy*. This interference between understanding language about action and performing a real action with our bodies suggests that while processing language, we perform motor imagery, using neural structures dedicated to motor control.

A third method, used by Stanfield & Zwaan (2001) and Zwaan et al. (2002), investigates the nature of visual object representations during language understanding. Zwaan and colleagues have shown that the implied orientations of objects in sentences (like *The man hammered the nail into the floor* versus *The man hammered the nail into the wall*) affect how long it takes subjects to decide whether images of those objects (such as a nail) was mentioned in the sentence. When the image of an object is seen in the same orientation as it was implied to have in the scenario described in the sentence (e.g. when the nail was described as having been hammered into the floor and was depicted as pointing downwards), it took subjects less time to perform the task than when it was in a different orientation (e.g. horizontal). The same result was found when subjects were just asked to name the object depicted. Zwaan and colleagues also found that when sentences implied that an object would have different shapes (e.g. an eagle in flight versus at rest), subjects once again responded more quickly to images of that object that were coherent with the sentence - having the same shape as they had in the sentence.

These convergent results from various methodologies suggest a concrete role for embodied perceptual and motor experiences in language understanding. Language understanders automatically mentally imagine, or *simulate* scenarios described by language. The mental simulations they perform can include motor detail at least to the level of the particular effector that would be used to perform the described actions, and perceptual information about the trajectory of motion (towards or away from the understander; up or down), as well as the shape and orientation of described objects. The neural imaging studies cited suggest that these simulations are executed by the very brain mechanisms responsible for perceiving the same percepts or performing the same actions.

These findings are not particularly surprising, given that motor and perceptual systems are known more generally to participate in higher cognitive functions like memory and imagery. A proliferation of recent neural imaging studies complements existing behavioral evidence that recalling motor experiences recruits cognitive mechanisms responsible for performing the same motor actions, by activating the same parts of the brain's motor system, just as recalling perceptual experiences, both in the visual and auditory domains, makes use of perceptual modality-specific neurocognitive structures (Barsalou 1999, Wheeler et al. 2000, Nyberg et al. 2001). Similarly, mental imagery involving motor control or visual or auditory perception yields activation of appropriate motor or perceptual brain areas (Porro et al. 1996, Lotze et al. 1999, Kosslyn et al 2001, Ehrsson et al. 2003). It thus seems that recalling, imagining, or understanding language about actions and percepts recruits brain structures responsible for performing the actions or perceiving the percepts that appear in the mind's eye.

Beyond this neuroscientific and psycholinguistic evidence, there are also good theoretical reasons to think that meaning involves the internal recreation of embodied experiences. For one, when considered as a real theory of actual language understanding, the abstract amodal symbol-manipulation view of meaning doesn't actually explain what meaning is or how it arises. Its shortcomings are clear from Searle's (1980) classic Chinese Room argument. In this thought experiment, a person is placed in a room, with a stack of Chinese characters and an instruction book. The book indicates, for any set of characters that appear through a slot into the room, what the appropriate characters to slide out are. Given an appropriate instruction book, this room and its contents may have the appearance to an outside observer of understanding and being able to produce coherent Chinese. However, the argument goes, with no more than a set of input-output rules expressed as arbitrary symbols, neither the room as a whole nor the person inside it actually understands Chinese. Abstract, amodal symbol manipulation doesn't provide an account of meaning as it is instantiated in humans, nor can it allow for the acquisition of meaning - only abstract rules can be learned in such a system.

By contrast, an embodied approach, in which understanding uses the internal execution of imagery qualitatively similar to the past experiences it is (re)created from, does offer an account of what meaning is and how it could be acquired. A simulation executed in response to language can be learned on the basis of previous associations between experiences of the world, including perceptual and motor knowledge, and the language used to describe them. As such, understanding language is qualitatively similar in some ways to experiencing the scenarios it describes. It is important to note that motor and perceptual experiences hold a privileged position in the study of mental simulation, only because their basic mechanisms and neural substrates are relatively well understood. Other dimensions of experience are also relevant to simulation; anything that is experienced, including affect, social interactions, subjective judgments, and other imagined scenarios can in principle be recruited to form part of a simulation.

As a sidenote, it may be important to consider the limits of previous experience in defining simulation. If meaning, as argued here, involves the activation of motor and perceptual (and other) representations of past experiences, how can language with counterfactual, or previously unexperienced meanings be understood? After all, one of the "design features" of human language is the possibility of describing things that do not exist (Hockett 1960), like the Easter Bunny or the current King of France. Moreover, because language is so important in helping children (and adults) learn about the world, it cannot be the case that linguistic meaning simply associatively reflects past experiences - if this were the case, then we could never learn anything new through language. However, a mental simulation-based account of meaning does not imply a purely behaviorist or empiricist perspective. In fact, there is good reason to believe that "mental images need not result simply from the recall of previously perceived objects or events; they can also be created by combining and modifying stored perceptual information in novel ways." (Kosslyn et al. 2001:635). Mental simulation involves the active construction by the conceiver of novel perceptual, motor, and affective experiences, on the basis of previous percepts, actions, and feelings, and while it is constrained and informed by these experiences, combinatorial and other creative capacities allow departures from them.

It thus seems that recalling, imagining, and most importantly for our purposes, understanding language about the world all use the perceptual, motor, and perhaps affective content of past embodied experiences to create a new, qualitatively similar, mental experience. But we have not yet seen how a theory of meaning can be constructed on the basis of these findings.

3. Simulation semantics

Understanding a piece of language can entail performing mental perceptual and motor simulations of its content. This implies that the meanings of words and of their grammatical configurations are precisely the contributions those linguistic elements make to the construction of mental simulations. The study of how different aspects of language contribute to the construction of mental imagery, and the corresponding theory of linguistic meaning as linguistic specifications of what and how to simulate in response to language, is known as *simulation semantics* (Bergen et al. 2003, Bergen and

Chang 2004, Bergen et al. 2004, Feldman and Narayanan 2004; see also foundational work on simulation and language by Bailey 1997 and Narayanan 1997).

A simulation-semantic view of meaning in language understanding involves three main components. The first is the activation and combination of parameterized representations of the simulative content that linguistic constructions evoke, using the words' meaning parameters as well as constraints provided by the grammar of the utterance. This step is known as the construction of a *semantic specification* - a parameterized set of instructions for subsequent simulation, and is discussed below. The second step involves the subsequent performance of a dynamic mental simulation, on the basis of the semantic specification, as described in the previous section. Finally, performing the simulation results in a set of mental experiences for the language understander, who has now internally perceived or acted out aspects of the content of the utterance. Knowledge of these experiences is propagated through the knowledge system according to what the understander believes he is supposed to do with them (examples might be enacting the content of the simulation, in the case of imperatives, storing the results for continued discourse, or updating beliefs about the speaker's (or someone else's) beliefs). The propagation of these inferences is discussed below.

To evaluate how pieces of language evoke simulation, we must first identify what these pieces of language are. It goes without saying that words of certain types, particularly content words such as nouns and verbs, contribute directly to the content of simulation, as shown by the studies descibed in Section 2, above. Less obviously, however, other linguistic elements, like function words and phrasal patterns have straightforward repercussions on simulation (Bergen and Chang 2004). Prepositions, for example, can indicate paths of motion to be included in a simulation, as in *The monkey rolled <u>down</u> the slide* or spatial relations as in *Each deaf child was seated <u>next to</u> a hearing adult*, among other things. Phrasal patterns can similarly contribute to the simulation. An example is the ditransitive, a clausal pattern with subject, verb, and two direct objects, exemplified by sentences like *John tossed me the book* or *I handed her the phone*. This sentence pattern can convey simulative content; in particular, it encodes the intended transfer of the second object (*the book*) from the subject (*John*) to the first object (*me*) (Goldberg 1995). Evidence for the meaningfulness of clausal patterns like the ditransitive comes from the possibility of using them to describe transfer scenarios, using verbs that do not themselves encode transfer (Kaschak and Glenberg 2000), such as *After nearly tripping over it, the injured woman turned and <u>crutched us the ball</u>*. Each chunk of language that pairs some form with some meaning, including morphemes, words, phrasal patterns, and even idioms, can be referred to as a *construction*, following the tradition in Construction Grammar (Goldberg 1995, Kay and Fillmore 1999).

How, then, do linguistic constructions evoke simulation? The evidence that processing utterances results in the activation of perceptual and motor imagery does not necessarily imply that it does so directly, without mediation. Instead, from the study of grammar, we find several types of evidence for an intermediary stage of representation between the form of linguistic constructions and the simulations they evoke. First, assembling words into an acceptable configuration during language production, or determining what the set of words in an utterance and their relations are during language understanding, does not require access to the detailed, encyclopedic meanings of the constructions involved. Rather, words are combined in part on the basis of generalizations about their meaning. Thus in order to participate as a verb in the ditransitive construction, a verb must be interpretable as encoding one of a number of types of intended successful transfer of one object to a recipient, or metaphorical types of transfer (such as communication, as in *She told him a story*). (Gropen et al 1989, Goldberg 1995). However, to be interpretable as a ditransitive sentence, no detailed semantic constraints are placed on the constituent verb at the level of motor or perceptual detail - we don't have to know what effector is used in what way, or what the direction or manner of encoded motion is. Thus, it isn't the case that a verb of hand-caused transfer like *toss* can occur in ditransitive constructions, while a verb of foot-caused transfer, like *kick* cannot. This ambivalence that constructions to be combined demonstrate towards simulative detail is well motivated. It would be quite burdensome to the language user if, upon processing any utterance, he were required to access the detailed perceptual and motor content of each of the possible (partial) interpretations of the utterance. Rather, linguistic units seem to encode only generalizations over aspects of the perceptual and motor content they trigger in simulation. All that is relevant about different types of action, for the ditransitive, is whether or not they can be interpreted as effecting a transfer of possession of one of various types. Such generalizations, also known as schematizations or parameterizations, constitute the purely linguistic representation of meaning, by being general enough to allow for operations of combination of linguistic elements. However, they facilitate simulation by being tightly linked to the simulative details they are schematized over.

A second piece of evidence that linguistic representations of meaning are distinct from the perceptual and motor content that they connect to derives from the possibility of understanding sentences that involve novel simulation content. If sentence understanding were strongly constrained by the actual motor and perceptual possibilities afforded by our real world experience, then we shouldn't be able to understand language that conflicts significantly with this experience. But we know that people can do this - sentences like *Opening up his mouth as wide as he could, the man ate first a table and chairs, then a bicycle, an umbrella, then a fish tank, and finished it off with a Volkswagon* are entirely interpretable and simulatable, even though it's quite clear that, due to their size and consistency, none of the object described as being consumed are usually edible.

Notice that it isn't just any word or set of words though that can combine with any other - if we replace *table and chairs* with *running of the Honolulu marathon* or *state of affairs*, the sentence above becomes uninterpretable. The verb *eat* seems to combine with direct objects that are physical objects, metaphorical objects like acts of communication (*he ate his words*) and several other classes, but the detailed affordances of whether an object can be eaten or not are not relevant until simulation. The point here is again that even if understanding language entails performing mental simulation, the act of preparing to perform a simulation cannot itself be directly contingent on existing detailed perceptual and motor knowledge, nor is it based exclusively on purely abstract criteria, like being a noun phrase. Rather, linguistic constructions like words and clausal patterns combine on the basis of generalized characteristics of the simulations that the words evoke.

If the meanings of linguistic constructions are parameterized representations of aspects of simulation, then comprehending an utterance involves fitting together the meanings of the constructions that make up that utterance. When understanding language, the comprehender must come up with a set or sets of constructions that fit together and best account for the utterance to be understood. A crucial part of this process, know as constructional analysis (Bryant ms.), is to bind together the meanings of the words, morphemes, phrasal constructions and idioms that make up the sentence. These must fit together not only in a formal arrangement (e.g. words are uttered in a particular order, with particular phonology), but also in a semantic arrangement (e.g. a verb's semantic subcategorization requirements are satisfied by elements representing the agent, patient, and so on). The result of this analysis process is a *semantic specification* defining a simulation to be run.

The simulation process, discussed in the previous section, is the second, and central, step in language understanding, on a simulation semantics approach. The third and final step is the propagation of inferences generated by the simulation (Narayanan 1997). Aside from creating an internal experience of the scenarios described by language, a major benefit of simulation is that it provides rich perceptual, motor, and affective detail about the content of an utterance, which the understander can then use to update her beliefs about the world. For example, when presented with a claim like *As a child, I ate nearly 50 hard-boiled eggs in an hour*, a hearer will run a mental simulation of the scenario and then update her knowledge about the speaker. Assuming that the speaker is to be believed, various inferences can be naturally made about the event. For example, the hearer is likely to believe that the speaker intended to eat 50 eggs but failed due to the discomfort caused by the endeavor. The discomfort probably involved stomach pain, accompanied by a dry mouth. The eater probably had to stuff the eggs into his mouth at times, and most likely felt rushed, since 50 eggs in an hour works out to one nearly every minute. Notice that this unique set of motor, perceptual, and affective inferences is absent from other, nearly identical sentences, like *As a child, I ate nearly 50 ants in an hour*. This suggests that the inferences are due to the construction of a specific model of what it would be like to perform or observe the described action. The understander's knowledge about this (purported) event in the speaker's childhood can be updated on the basis of simulation to facilitate future behavior, including future language understanding. For example, if either of the sentences above (about eating eggs or ants) is followed by an utterance like *I was in the hospital for a week*, the hearer immediately understands the implied causal link between the two events. Because she already believes the speaker to have completely stuffed his stomach through egg eating, or to have put a large number of insects in his body, the understander easily sees the causal relationship between the two predications.

Note that inference propagation could in principle proceed without first engaging the simulation mechanism. Through associative learning, a language user could come to know that whenever a certain piece of language is uttered, she should simply update her beliefs in some particular way. For example, pervasive interjections like *sure* and *yeah* can in many cases be seen as having a direct path, or shortcut, to belief-updating processes, with no need to wait for simulation. Recognizing simulation as an important component of language understanding does not necessarily entail that simulation is always necessary for understanding every piece of every utterance, although, as argued here and elsewhere, the simulation process appears to be at work in a great deal of language understanding.

By contrast with the classical theory of word meaning described above, the approach described here views meaning as not based on abstract, amodal symbols, but rather on schematizations over aspects of simulation, in combination with the simulation itself that those schematizations generalize over, and the subsequent propagation of inferences. Rather than being simply static representations, word meanings combine static, parameterized knowledge with the dynamic execution of the simulations they specify. On this view, meaning can be conveyed by words and morphemes in very much the same way that it can be for larger units, like idioms and sentential constructions.

4. Figurative language and mental simulation

Meaning has been argued above to result from aligned semantic parameterizations producing a dynamic, modal simulation, with the inferences from that simulation subsequently propagated through the conceptual system. Various scholars have argued, pursuant to the standard pragmatic model, that figurative language makes use of different mechanisms than those responsible for literal language processing. For instance, Giora et al. (1998) argue that increased reading time for sentences in figurative contexts versus literal contexts results from an additional pragmatic processing component tacked on to the end of literal processing. Such claims, however, are based on the notion of literal language as static, lexical, and compositional. By contrast, a simulation semantics account of literal meaning, due to its necessary incorporation of dynamic simulation and inference propagation, straightforwardly provides a framework within which figurative language processing makes use of the very same mechanisms that literal language processing does. What is advantageous about such an account, beyond its greater parsimony, is that it permits the respective details of figurative and literal language understanding processes, as well as their interactions with other factors in processing like conventionality and idiomaticity, to be directly compared. It also allows the advantages of a simulation-based account of literal language understanding, like the creation of dynamic, embodied inferences, to be extended to the study of figurative language.

Figurative language understanding could naturally differ from or be similar to literal language processing in any of the three components of simulation semantics (parameterization, simulation, or inference propagation) or the relations among them. We argue below that meaningful language, whether figurative or literal, makes use of the same set of understanding mechanisms. Moreover, the apparent figurativeness of certain types of language results from three main factors: idiomaticity, conventionality, and indirect relations between simulation and the inferences that are propagated on the basis of it. All three of these factors are straightforwardly dealt as characteristics of or relations among the uses of components of a simulation-based language understanding system. While figurative language is often treated as a unified phenomenon, the close inspection of figurative language processing below indicates that, as pointed out by Gibbs (1994), the various different figures (metaphor, metonymy, idiom, irony, hyperbole, etc.) may have very different bases.

Dimensions of figurativeness

There are at least three factors that conspire to create the appearance of figurativeness. One, discussed above, is the degree to which the semantic specification is affected by idioms and other larger-than-lexical constructions. The same semantic specifications can be constructed on the basis of language that includes or does not include idiomatic constructions, but language is more likely to appear figurative when an idiom is present. Since idioms by definition do not have compositional meaning, and since language that uses them usually has an alternative, compositional interpretation, idiomatic language displays an apparent mismatch between the "literal" meaning of the utterance and the content of the semantic specification, and thus, the simulation as well. But while idiomaticity contributes to figurativeness, it isn't uniquely constitutive of it.

A second dimension of figurativeness involves the degree to which a particular contribution to the semantic specification is conventionally associated with a linguistic construction. Some lexical items, like *see*, have metaphorical senses (in this case, 'to understand') which are so strongly conventionalized that they must be encoded as part of the parameterized semantic encoding of the words. Other words can be used figuratively in relatively novel ways, like *microscope* in *The judges put my rendition of* Ave Maria *under a microscope*. Since *microscope* has no conventionalized meaning pertaining to thorough non-visual inspection, it would not produce a parameterized representation of such to be fed into simulation. Instead, as we see below, novel metaphorical language, like this use of *microscope,* may provide content for a source-domain simulation (in this case involving the domain of visual inspection), whose content is subsequently used for inferences about the target domain (in this case, evaluation). Again, though, while novelty contributes to figurativeness, it does not do so alone.

Finally, the relation between the simulation that is performed and the inferences that are propagated may contribute to figurativeness. What is usually seen as literal language is language in which inferencing processes may proceed directly on the basis of the simulation of the sentence. Thus, the sentence *The students were glued to their seats*, when interpreted literally, yields inferences that take the content of the utterance as entirely factual. Of course, there are many uses of literal language in which simulation should not yield a direct updating of one's beliefs about the world. For example, since speakers are not equally reliable, determining whether an understander should update her beliefs about the world or her beliefs about the speaker on the basis of a simulation may depend on the speaker. If your student-gluing colleague happens to be a known pathological liar, you are much more likely to infer, assuming the literal interpretation of the sentence, that the simulation you have constructed is to be attributed to his claims, rather than to your beliefs about the real world.

The study of mental spaces (Fauconnier 1985) treats attributions of meaning content to different people's beliefs, hypothetical worlds, and so on in great detail. However, figurative language tends to differ from literal language in how simulation relates to inference. As noted by Gibbs (1994), there are many varieties of figurative language understanding, including metaphor, metonymy, irony, hyperbole, and idiom, and these different types of processing may be underlain by a diverse set of relations between simulation and inference.

Simulation in different types of figurative language

The various types of figurative language differ in the ways that simulation may relate to subsequent inference. Metaphorical language has been the object of more careful scrutiny than other types figurative language use, so we begin with it.

Consider a sentence like *The judges put my rendition of* Ave Maria *under a microscope*, which makes use of a conventional conceptual metaphor, UNDERSTANDING IS SEEING (Lakoff and Johnson 1980) but no idiom or conventional metaphorical word meaning. Understanding a sentence using metaphorical language like this has been argued (Narayanan 1997) to depend on first performing a mental simulation that includes elements of the concrete source domain (visual inspection), followed by inferences being propagated about the target domain of thought (about my musical performance). In particular, the inferences would include the observation that the consideration of by performance was extremely detailed. There is some psychological evidence of access to a metaphorical target domain during metaphorical language processing. Gibbs et al (1997), for example, argue that while processing idiomatic metaphorical expressions (like *He blew his stack*), readers automatically and unconsciously activate the metaphorical source domain (*heat*), as demonstrated by the priming of source domain language. Similarly, Boroditsky (2000, 2001) has shown that conceptual metaphors in which time passing is understood as (vertical or horizontal) motion through space lead to priming effects in which performing spatial reasoning influences temporal reasoning, but not the reverse. Results like these may imply that thinking about target domain concepts (in our language, performing simulations involving them) involves activating source domain imagery.

While simulation semantics predicts a role for source-domain simulation in much of metaphorical language understanding, the nature of this simulation is still very much at issue. One suggestion that comes out of research on conceptual integration (Fauconnier and Turner 2002) is that, just like the language that produces it, metaphorical simulation incorporates elements of both the source and the target domain (Grady et al. 1999). For example, in the judges-and-microscope sentence, as in most metaphorical language, elements of both source and target are mentioned, and it could be that a single simulation is constructed, in which the judges are using a microscope to visually inspect something, which is a miniature me, performing my Ave Maria. If this is the case, it can explain where the inferences that the judges attended to minutiae in my performance comes from - a hyper-literal simulation.

What is usually referred to as metaphorical language, like all the other sorts of figurative language discussed in this section, can be more or less idiomatic, and more or less novel. As a result, the nature of the simulation used in the processing of this figurative language and its relation to inference may range from extremely conventional to extremely figurative (this point is argued for metonymy, discussed below, by Papafragou (1996)). This is a point we will return to below.

A different sort of figurative language, metonymy, can also be seen as resulting in a disjunct between the simulation and resulting inferences. In most metonymies, a word or set of words that identify a referent (the *trigger*) are used to identify a second referent (the *target*) that is pragmatically related in some way to that first referent (Fauconnier 1985). It's important to point out that on some accounts (e.g. Nunberg 1995), metonymy is a strictly linguistic phenomenon - one in which a particular word or class of words are conventionally used to refer to referents or predications that their usual (or literal) referents or predicates are related to in some way. On such theories in which metonymy is a matter of words, one word simply stands in for another, such that the semantic specification specifies Madonna's albums as the thing being kept somewhere, and the word *Madonna* only serves to trigger the metonymy whereby this information is placed there. Thus, on this view, metonymy involves the indirect construction of a semantic specification, either through conventional meanings associated with the particular words (like the various meanings of *White House*) or through a conventional linguistic mechanisms of construal whereby words of certain classes are automatically assigned additional, metonymic meanings.

While there certainly are conventional metonymies, in many cases reference in metonymies seems to be based on more general associative and encyclopedic cognitive operations (Lakoff 1987, Panther and Radden 1999, Panther and Thornburg 2003). In these cases of metonymy, reference resolution is too complex or novel for pre-existing linguistic devices to suffice. Instead, only through accessing knowledge about the trigger and the scenario being described - that is, through simulation - can the target be determined. For example, an utterance like *Let's hope this election year we don't see another Florida* is metonymic, since *Florida* usually refers to a place, and not an election scandal. But it would be hard to argue either that the word *Florida* has a referent which is the particular election scandal that took place there in the year 2000, or that there is a purely <u>linguistic</u> process whereby names for locations where disasters took place can be used to refer to the disasters themselves. Much more likely is the possibility that reference resolution happens on the basis of general conceptual knowledge about elections and Florida, invoked by the semantic specification.

It has been argued that metonymic language of this second type, where conventional linguistic knowledge doesn't satisfactorily account for the trigger-target relation, is understood on the basis of an integrated simulation (Alac and Coulson 2002), as described above for metaphor. In such a simulation, the trigger of the metonymy (e.g. *Madonna*) is blended together with target-oriented description of the location of an object, such that the simulation includes some combination of a real Madonna with artistic creations by her or a representation of her under my sister's pillow. As with metaphor, on this account, inferences are subsequently propagated not about this blended simulation, but rather about one in which it is the resolved, intended target referent that is located under my brother's pillow.

A third, well-studied type of figurative language, *fictive motion* (Talmy 1996), uses motion language to describe static scenes, as in sentences like *The fence <u>runs</u> from the house down to the river* or *Highway 1 <u>weaves</u> around the island.* In a set of language comprehension studies, Matlock (To Appear) demonstrated that the time subjects took to understand fictive motion sentences was influenced by how quickly one could move along the described paths. For example, a sentence like *The path followed the creek* was processed faster when it followed a paragraph describing an athletic young man who jogs along the path than when it followed one describing an old man who had difficulty walking all the way down the path. Similarly, characteristics of the path itself like its distance or difficulty to navigate, were found to influence processing time in the same direction - the longer it would take the mover to travel the path, the longer it took subjects to process the fictive motion sentence. Even if the precise relation between fictive and literal motion, as well as the content of simulation in response to fictive motion remain to be clarified, this work nevertheless implies that fictive motion, like literal motion, is processed using mental simulation of described scenarios.

Other sorts of figurative language, like hyperbole and irony, have similarly been hypothesized to involve the execution of a simulation that differs from the inferences that are subsequently propagated, in different ways. Hyperbole, when the process isn't shortcut by lexical or idiomatic specifications of an intended figurative simulation, yields simulation that is false and potentially impossible in terms of its scale (Bergen and Binsted 2004, To Appear). For example, *I laughed so hard I burst my spleen* evokes a simulation in which the speaker's spleen bursts from heaving laughter, but the inferences that are subsequently propagated do not lead the understander to conclude that the speaker is in need of immediate medical treatment. Rather, hyperbole yields scaled down inferences. Irony is unique among the forms of figurative language understanding discussed here, since its literal and figurative meanings are completely incompatible. Aside from conventionalized ironic language (*Yeah, right!)*, and contextually expected ironic language can be hypothesized, like negation, to yield a simulation which encodes the opposite of the inferences to be propagated (Hasson and Glucksberg Ms., Giora et al. To Appear). Irony differs from negation in that there is often no overt lexical or morphological marker in the case of irony that the inferences to be propagated should be the opposite of the simulation, while negation explicitly marks this morphologically.

All of the routes to figurative language understanding described here require language users to be able to distinguish between the content of a simulation and the inferences that should be propagated on the basis of it. The exact relationship between these is often specified by particular linguistic constructions, such as the use of constructions like X is Y, which can trigger a metaphorical interpretation (*My roommate is a wild boar*) or X is so Y that Z, which can lead to a hyperbolic interpretation (*It was so cold in New York that the strippers in central park were just describing themselves.*). Another example of a type of cue provided by a linguistic construction is the intonational indicators of irony. Linguistic incongruities are also triggers for alternative processes of inference propagation, as with metonymy and sometimes metaphor, or much of idiom. For example, in *Congress once again sidestepped The Equal Rights Amendment*, the direct object is not a physical object, and thus does not match even the broad semantic requirements of the verb *sidestep*. Indirect relationships between simulation and inference propagation can also arise from mismatches between the understander's current beliefs and the content of a simulation. If an understander believes that the current discussion is about the value of stocks and not their physical height, then running a simulation of a sentence like *Martha Stewart's stocks inched up minutes before the verdict was announced*, in which there is a change in physical height, leads to inferences that are incompatible with the understander's current beliefs. As such, the inferences to be propagated are modified appropriately, using the existing metaphorical mapping between quantity and height from the conceptual metaphor MORE IS UP (Lakoff and Johnson 1980).

Degrees of figurativeness and simulation

These proposed characterizations of the different varieties of figurative language processing - involving disjunctions between conventional constructional parameterizations and simulation, or between simulation and inference propagation, interact with the other two dimensions of figurativeness discussed above. Language analyzed as making use of a conceptual metaphor, for example, may be more or less conventionalized and idiomatic, and as a result may be more or less likely to produce source domain or blended simulation. If figurative meanings are part of the conventionalized meaning pole of the words in questions, very little or no source domain simulation may be necessary, and these words might behave instead just like literal language about the same target domain. For example, since *see* has a conventionalized meaning, meaning something like 'understand', it may well have among its possible contributions to the semantic specification something that looks just like the contribution made by the verb *understand*. Similarly, idiomaticity may interact with the mechanisms described above, like metaphor and irony. Literal idiomatic language (like *What's that scratch doing on the table?*) encodes a direct connection to its particular imagery. Figurative idiomatic language (like *Don't have a cow!*) may similarly encode a direct relation to target domain simulation, thus shortcutting the step of performing source domain-based imagery.

There is however, some evidence that even when metaphorical idioms are entirely conventionalized, source domain details are still activated (Gibbs et al 1990, 1997). This finding opens up the possibility that processing figurative language may sometimes entail performing multiple simulations, constituting for example a compositional and an idiomatic interpretation of the same metaphorical sentence. Although it isn't usually interpreted in this way, some psycholinguistic research showing that figurative language takes longer to interpret than literal language (e.g. Temple and Honeck 1999) be an indication of multiple simulations. However, this line of reasoning assumes that simulations must necessarily proceed in sequence. It is also possible that certain types of simulation could be performed in parallel. Simulation makes use of the same neural structures dedicated to the perception or action, and it is known that the more incompatible two actions or percepts are, the more difficult they are to perform or perceive simultaneously. Similarly, as shown by the behavioral evidence for simulation described above, incompatibilities between actions or percepts and mental simulations yield processing difficulties (Glenberg and Kaschak 2002, Bergen et al. 2003, Richardson et al 2003). It is thus unlikely that incompatible simulations could be processed simultaneously. However, compatible simulations, i.e. those that activate the same neural structures in the same way, or those that activate unrelated neural structures, might be capable of running efficiently in parallel.

Another way in which figurative language may vary is in terms of the degree of simulation required to understand it. As discussed above, some language, including literal language, is so closely associated with the inferences that can be made on the basis of it - for example *Yeah, right!* implies ironically that the speaker doesn't believe a particular proposition - that it may partially or completely skip simulation. Such shortcuts from semantic specification to inferencing are entirely predictable on the basis of associative learning; when a construction co-occurs frequently and with a high cue validity with a pattern of inferences, links may arise that directly connect the two, short-cutting simulation. Figurative language displays evidence of this behavior. For example, metaphorical expressions like *See what I mean?* or *You follow?* entail that the speaker wants confirmation that the hearer is understanding the discourse, but because of their frequent use, associated with this particular inference, little or no simulation may be required for appropriate inferences to be drawn and beliefs to be subsequently updated.

5. Conclusions

Figurative language understanding has been argued here to rely on the same mechanisms that are used for literal language understanding. Evidence from neural imaging and behavioral experimentation shows that mental imagery has an important role in language understanding. We have argued that the process of internal simulation interfaces on one side with a constructional analysis process that aligns parameterized representations of simulation detail, encoded on linguistic constructions. On the other side, it interacts with inference propagation mechanisms. Rather than offering a unified account of all figurative language, it has been suggested here that figurative language can diverge from literal language in a number of ways, through idiomaticity, novelty, or through indirect relations between the inferences generated through simulation and those subsequently propagated.

If simulation really is a vital component of figurative language understanding, then we should see effects of simulation in language processing studies. For example, simulation implies as mentioned above that language understanding processes that require multiple incompatible simulations should take longer to process than those that require multiple, compatible simulations. Processing any utterance, literal or figurative, also requires the propagation of inferences on the basis of simulation, whether the utterance is interpreted literally or figuratively, as described above. It would therefore not be expected that figurative language should not be necessarily processed any more slowly than literal language, unless it involves the performance of multiple simulations. Since conventionality of the association between language and inferences, as well, perhaps, as idiomaticity, determines whether and to what extent a simulation is run, this should be the main factor in determining time to process a given piece of language literally or figuratively.

Similarly, we should expect to find measurable effects of simulations even when the inferences propagated are not directly based on the simulation. For example, when people process sentences that make use of perceptual or motor imagery to metaphorically describe domains that are more abstract, we should find interference effects like the ones seen in the behavioral experiments described in Section 2 above, though perhaps on a relatively short timescale.

To conclude, a simulation semantics perspective on figurative and literal language provides clarification of the age-old insight that figurative language is language differs from but is based on literal meaning. In some cases, linguistic knowledge and processes (including idiom) can yield multiple relations between words and semantic specifications, some of which will be more compositional and concrete. In other cases, and in some overlapping ones, the content of simulation can produce inferences that differ from the inferences that are subsequently propagated through the understander's belief system. Either way, the linguistic form of figurative language is disconnected from what the understander learns about the world on the basis of it. In an embodied approach to language, the experiences an individual has had in the world are viewed as vital to the architecture of their cognitive faculties and their behavior. It seems that internal simulation, based on previous action and perception in the world, is essential to understanding literal and figurative language processing.

References

Alac, Morana and Seana Coulson. 2002. The man, the key, or the car: Who or what is parked out back? Paper presented at CSDL-6, Houston.

Bailey, D. (1997). When push comes to shove: A computational model of the role of motor control in the acquisition of action verbs. Ph.D. Dissertation, Computer Science Division, University of California, Berkeley

Barsalou, L. W. 1999. Perceptual symbol systems. Behavioral and Brain Sciences, 22, 577-609.

Bergen, Benjamin and Kim Binsted. 2004. The Cognitive Linguistics of Scalar Humor. In Michel Achard and Suzanne Kemmer (Eds.) Language, Culture, and Mind. CSLI.

Bergen, Benjamin and Kim Binsted. To Appear. Embodied Grammar and Humor. Humor.

Bergen, Benjamin and Nancy Chang. 2004. Embodied Construction Grammar in Simulation-Based Language Understanding. In Jan-Ola Östman and Miriam Fried (Eds.), Construction Grammar(s): Cognitive and Cross-language dimensions.

Bergen, Benjamin, Nancy Chang, and Shweta Narayan. 2004. Simulated Action in an Embodied Construction Grammar. Proceedings CogSci.

Bergen, Benjamin, Teenie Matlock, and Srini Narayanan. Ms. Mental simulation in concrete, abstract, and metaphorical language processing.

Bergen, Benjamin, Shweta Narayan, and Jerome Feldman. 2003. Embodied verbal semantics: evidence from an image-verb matching task. Proceedings of the Twenty-Fifth Annual Conference of the Cognitive Science Society.

Boroditsky, Lera. (2000). Metaphoric Structuring: Understanding time through spatial metaphors. Cognition, 75(1), 1-28.

Boroditsky, L. (2001). Does language shape thought? English and Mandarin speakers' conceptions of time. Cognitive Psychology, 43(1), 1-22.

Bryant, Johno. 2003. Constructional Analysis. UC Berkeley M.S. Report.
Dascal, Marcelo, 1987. Defending literal meaning. Cognitive Science 11: 259–281.
Ehrsson, H.H., Geyer, S., and Naito, E. 2003. Imagery of voluntary movement of fingers, toes, and tongue activates corresponding body-part specific motor representations. J. Neurophysiol. 90: 3304-3316.
Fauconnier, Gilles. 1985. Mental Spaces. Cambridge, Mass.: MIT Press.
Fauconnier, Gilles and Mark Turner. 2002. The Way We Think. New York: Basic Books.
Feldman, Jerome and Srinivas Narayanan. 2004. Embodied Meaning in a Neural Theory of Language. Brain and Language.
Gibbs, Raymond. 1994. The Poetics of Mind. Figurative Thought, Language and Understanding, Cambridge University Press, Cambridge.
Gibbs, Raymond. 1998. The fight over metaphor in thought and language, in: A.N. Katz, C. Cacciari, R.W. Gibbs, M. Turner (Eds.), Figurative Language and Thought, Oxford University Press, Oxford, 1998, pp. 88–118.
Gibbs, Raymond W. and O'Brien, Jennifer E. (1990), "Idioms and Mental Imagery: The Metaphorical Motivation for Idiomatic Meaning", Cognition, 36:1, 35-68.
Gibbs, R., Bogdonovich, J., Sykes, J., & Barr, D. (1997). Metaphor in idiom comprehension. "Journal of Memory and Language," 37, 141-154.
Giora, R., Balaban, N., Fein, O., & Alkabets, I. (To Appear). Negation as positivity in disguise. In: H. L. Colston, & A. Katz (Eds.), Figurative language comprehension: Social and cultural influences. Hillsdale, NJ: Erlbaum.
Giora, Rachel, Ofer Fein and Tamir Schwartz, 1998. Irony: Graded salience and indirect negation. Metaphor and Symbol 13: 83–101.
Glenberg, A. M. & Kaschak, M. P. 2002. Grounding language in action. Psychonomic Bulletin & Review.
Goldberg, A. 1995. Constructions: A Construction Grammar Approach to Argument Structure. Chicago: U Chicago Press.
Grady, J., Oakley, T., & Coulson, S. (1999). Conceptual blending and metaphor. In R. Gibbs (Ed.) Metaphor in Cognitive Linguistics. Amsterdam: John Benjamins, pp. 101-124.
Gropen, J., Pinker, S., Hollander, M., Goldberg, R. & Wilson, R. (1989) The learnability and acquisition of the dative alternation in English. Language, 65, 203-257.
Hasson, Uri and Sam Glucksberg. Ms. Does understanding negation entail affirmation?
Hauk, O., Johnsrude, I. & Pulvermüller, F. 2004. . Somatotopic representation of action words in human motor and premotor cortex. Neuron. 41(2): 301-7.
Hobbs, J.R., M.E. Stickel, D.E. Appelt, P. Martin. 1993. Interpretation as abduction, Artificial Intelligence 63:69–142.
Hockett, C.F. 1960. The origin of speech. Scientific American, 203, 88-96.
Kaschak, M. P., & Glenberg, A. M. (2000). Constructing meaning: The role of affordances and grammatical constructions in sentence comprehension. Journal of Memory and Language, 43, 508-529.
Katz, Jerrold J., 1977. Propositional structure and illocutionary force. New York: Thomas Y. Crowell.
Kay, Paul & Charles J. Fillmore. 1999. Grammatical constructions and linguistic generalizations: The What's X doing Y? construction. *Language* 75/1: 1-33.
Kosslyn, S.M., Ganis, G., and Thompson, W. L. (2001). Neural foundations of imagery. Nature Reviews Neurosci, 2, 635 -642.
Lakoff, George. 1987. Women, fire, and dangerous things. Chicago: University of Chicago Press.

Lakoff, George and Mark Johnson. 1980. Metaphors we live by. Chicago: University of Chicago Press.
Langacker, Ronald. 1991. Concept, Image, Symbol. Mouton.
Lotze, M., Montoya, P., Erb, M., Hülsmann, E., Flor, H., Klose, U., Birbaumer, N., & Grodd, W. (1999) Activation of cortical and cerebellar motor areas during executed and imagined hand movements: An fMRI study, Journal of Cognitive Neuroscience, 11(5): 491-501
Markert, Katja and Udo Hahn (1999): Grounding figurative language use in incompatible ontological categorizations. In *Proceeedings of the 21st Annual Conference of the Cognitive Science Society (CogSci '99)*, pp361-366. Vancouver, Canada.
Markert, Katja and Udo Hahn. 2002. Metonymies in Discourse. (ps) In Artificial Intelligence 135 (1-2), pp. 145-198.
Matlock, T. (To Appear). Fictive motion as cognitive simulation. Under review.
May, R. 1985. Logical Form. Cambridge: MIT Press
Michaelis, Laura. 2004. Type shifting in construction grammar: An integrated approach to aspectual coercion. Cognitive Linguistics 15-1: 1–67.
Narayanan, S. (1997). KARMA: Knowledge-based active representations for metaphor and aspect. Ph.D. Dissertation, Computer Science Division, University of California, Berkeley
Nunberg, Geoff. (1995), "Transfers of meaning", Journal of Semantics 12: 109-132.
Nyberg, L., Petersson, K.-M., Nilsson, L.-G., Sandblom, J., Åberg, C., & Ingvar, M. (2001). Reactivation of motor brain areas during explicit memory for actions. NeuroImage, 14, 521-528.
Ortony, Andrew. 1980. Some psycholinguistic aspects of metaphor, in: R. Honeck, R. Hoffman (Eds.), Cognition and Figurative Language, Erlbaum Associates, pp. 69–83.
Panther, Klaus-Uwe and Günter Radden (eds.). 1999. Metonymy in Language and Thought. John Benjamins.
Panther, Klaus-Uwe and Linda L. Thornburg (eds.). 2003, Metonymy and Pragmatic Inferencing. John Benjamins.
Papafragou, Anna. 1996. Figurative language and the semantics-pragmatics distinction. Language and Literature 5: 179-93.
Porro CA, Francescato MP, Cettolo V, Diamond ME, Baraldi P, Zuian C, Bazzocchi M, & di Prampero PE (1996) Primary motor and sensory cortex activation during motor performance and motor imagery: a functional magnetic resonance imaging study. J Neurosci 16:7688–7698.
Pulvermüller, F., Haerle, M., & Hummel, F. 2001. Walking or Talking?: Behavioral and Neurophysiological Correlates of Action Verb Processing. Brain and Language 78, 143–168.
Pustejovsky, J. 1991. The generative lexicon, Comput. Linguistics 17 (4): 409–441.
Richardson, D. C., Spivey, M. J., McRae, K., & Barsalou, L. W. 2003. Spatial representations activated during real-time comprehension of verbs. Cognitive Science.
Searle, John R., 1978. Literal meaning. Erkenntnis 13: 207–224.
Searle, John. 1980. "Minds, Brains, and Programs." Behavioral and Brain Sciences 3, 417-424.
Stanfield, R. A. & Zwaan, R. A. 2001. The effect of implied orientation derived from verbal context on picture recognition. Psychological Science, 12, 153-156.
Talmy, L. (1996). Fictive motion in language and "ception". In P. Bloom, M.A. Peterson, L. Nadel, & M.F. Garrett (Eds.), Language and space (pp. 211-276). Cambridge, MA: MIT Press.
Temple, John and Robert Honeck, 1999. Proverb comprehension: The primacy of literal meaning. Journal of Psycholinguistic Research 28: 41–70.
Tettamanti, M., Buccino, G., Saccuman, M. C., Gallese, V., Danna, M., Perani, D., Cappa, S. F., Fazio, F., & Rizzolatti, G. Unpublished Ms. Sentences describing actions activate visuomotor execution and observation systems.

Wheeler, M. E., Petersen, S. E., & Buckner, R. L. (2000). Memory's echo: Vivid remembering reactivates sensory specific cortex. Proc. Natl. Acad. Sci. USA 97: 11125–11129.

Zwaan, R. A., Stanfield, R.A., & Yaxley, R. H. 2002. Do language comprehenders routinely represent the shapes of objects? Psychological Science, 13, 168-171.