Chapter 15

Language comprehension is both embodied and symbolic

Max Louwerse and Patrick Jeuniaux

15.1 Introduction

Over the last decades a divide seems to have emerged in the cognitive sciences between embodied and symbolic approaches to language understanding. Embodied approaches argue that language comprehension requires activation of our experiences with the world, whereas symbolic approaches argue that language comprehension relies on interdependencies of words. The current paper argues that language comprehension is both embodied and symbolic. According to the *symbol interdependency hypothesis* comprehenders can ultimately ground symbols, but they also can rely on interdependencies across symbols as a shortcut to the meaning of words. An overview is given of the evidence supporting this hypothesis suggesting that embodied representations are activated under certain conditions and ultimately tend to be encoded in language structures.

Attempts to answer the question to what extent language comprehension is embodied or symbolic are not new. For instance, around 360 BC, Plato describes a discussion between Cratylus, Hermogenes, and Socrates, with Socrates asking:

> Then if you admit that primitive or first nouns are representations of things, is there any better way of framing representations than by assimilating them to the objects as much as you can; or do you prefer the notion of Hermogenes and of many others, who say that names are conventional, and have a meaning to those who have agreed about them, and who have previous knowledge of the things intended by them, and that convention is the only principle; and whether you abide by our present convention, or make a new and opposite one, according to which you call small great and great small – that, they would say, makes no difference, if you are only agreed. Which of these two notions do you prefer? (Plato 1902)

Over the last few decades the field of cognitive science seems to be divided between those who prefer one of Socrates's notions over the other. On the one hand, embodied accounts of language processing argue that symbols are fundamentally grounded in embodied experience (whereby the processing of symbols is modal), and indeed there is compelling evidence that symbols can be grounded, subconsciously or consciously (Barsalou 1999; Glenberg 1997; Harnad 1990; Pulvermüller 1999; Searle 1980; Zwaan 2004). For instance, we all know intuitively that we *can* imagine the colour red when reading the word 'rose', the smell when reading the word 'perfume', and the loud noise when reading the word 'thunder'. However, the specific stance of the embodied accounts of language processing, is that all words *must always* activate perceptual representations of their referents (whereby the processing of meaning is modal) in order to activate meaning.

On the other hand, symbolic accounts of language processing argue that symbols retrieve much of their meaning by their association with other symbols (whereby the processing of meaning is amodal). Indeed, there is also compelling evidence that symbols can derive meaning from other symbols (Fodor 1975; Kintsch 1998; Landauer and Dumais 1997; Pylyshyn 1984). For instance, we know that a rose is a flower because 'rose' appears in linguistic contexts similar to the contexts in which 'flower' appears. The same is true for associations between 'perfume' and 'smell' and for 'thunder' and 'loud'.

Though the discussion whether language comprehension is fundamentally symbolic or embodied has become a centuries-old debate (see Ogden and Richards 1923 for an overview), this debate recently received new impetus. For instance, Pecher and Zwaan's (2005) edited book *Grounding Cognition: The Role of Perception and Action in Memory, Language, and Thinking* leaves little room for symbolic representations while showing that language is always grounded (in modal representations), whereas Landauer *et al.*'s (2007) edited book *Handbook of Latent Semantic Analysis* pays little attention to embodied representations while showing the strengths of higher-order relationships between (amodal) language units.

Advocates of embodied versus symbolic accounts of language processing defend their side fiercely, as illustrated in the following quotes.

> … (amodal) propositions are 'a convenient shorthand' for representing information. Indeed, as these authors and the many researchers inspired by them have shown, impressive results can be obtained by using this shorthand, but as is shown here, it would be a mistake to elevate the shorthand to the status of longhand (Zwaan 2004, p. 56–7).
>
> Therefore, it is gratuitous to conclude that [latent semantic analysis (LSA)] is wrong in principle from the observation that it is sometimes wrong as implemented and trained ... The use of these [embodiment] ideas to oppose [hyperspace analogue to language (HAL)] and LSA is a case of what Dennett calls an 'intuition pump', pushing the introspective mystery of a mental phenomenon to discredit a mechanistic explanation (Landauer 2002, p. 66–7).

At the same time, however, neither side argues that aspects from the other side should be excluded from the theoretical framework.

> …Thus, it can be argued that the information captured by amodal propositions forms a subset of the information captured by a perceptual analysis (Zwaan 2004, p. 56).
>
> … [Embodiment] is an appealing idea; it offers the beginnings of a way to explain the relation between some important aspects of thought, language, and action that appear to capture analogue properties of the cognition of experience (Landauer 2002, p. 67).

It is noteworthy that the literature on the debate between embodiment and symbolism largely follows the distinction between psycholinguistic and computational linguistic approaches. The lion's share of embodiment publications is experimental in nature, ranging from reaction time to event related-potential experiments (Boroditsky 2000; Fincher-Kiefer 2001; Glenberg and Kaschak 2002; Wiemer-Hastings and Xu 2005; Zwaan *et al.* 2002). Far fewer computational models have been developed that integrate embodied representations (but see Cangelosi and Harnad 2000; Roy and Pentland 2002; Vigliocco *et al.* 2004; Westermann and Reck Miranda 2004).

At the same time, the majority of symbolism publications is computational in nature, showing how data are related to corpus linguistic findings (Burgess 1998; Howard and Kahanna 2002, Kintsch 1998a,b; Landauer 2002; Landauer and Dumais 1997; Lund and Burgess 1996; Steyvers *et al.* 2004). Fewer experimental studies have been conducted arguing that language processing is to a large extent amodal (but see Bransford and Franks 1971; Kintsch and Keenan 1973; Van Dijk and Kintsch 1983). Overall, there seems to be a methodological preference for one side or the other. An interdisciplinary account incorporating psycholinguistics and computational linguistics is probably as hard to find as is a unified account incorporating embodiment and symbolism.

15.2 **Language comprehension is embodied**

There is no doubt that linguistic symbols *can* be grounded in the physical world of the comprehender's experiences. A range of theoretical and experimental evidence shows that, in the process of language comprehension, the comprehender activates embodied representations. Philosophical arguments have posited that language comprehension *must* involve more than translating one symbol into another (Harnad 1990; Searle 1980). Whether or not these thought experiments prove their case remains open for discussion (Bishop and Preston 2001), but at least strong theoretical arguments are posed against a solely symbolic account of language comprehension.

Indeed, there is overwhelming experimental evidence for embodied representations. For instance, Pecher *et al.* (2003) and Spence *et al.* (2000) found that stimuli from the same modality resulted in faster processing times than stimuli from different modalities (even when the modalities were described in words), suggesting that modalities don't seem to get recoded into an amodal symbolic representation. Furthermore, there is evidence that motor movements of comprehenders match those described in the linguistic input. Klatzky *et al.* (1989) showed that the comprehension of verbally described actions was facilitated by preceding primes specifying the motor movement. Glenberg and Kaschak (2002) found an action-sentence compatibility effect (ACE) with motor imagery described in a sentence effecting motor actions. Zwaan *et al.* (2002) showed that response times for sentences were shorter with pictures matching a sentence than mismatching pictures. Zwaan and Yaxley (2003) showed that the spatial configuration of items presented on a screen affects semantic judgements about them, suggesting that visual representations get activated during language comprehension.

The argument that nonlinguistic representations are tightly coupled to language is also found in eye-tracking studies showing specific patterns of eye gaze influenced by text comprehension or problem solving activities of various degrees of complexity (Altman and Kamide 1999; Chambers *et al.* 2002; Demarais and Cohen 1998; Matlock and Richardson 2004; Spivey and Geng 2001; Spivey *et al.* 2000, 2002). A slightly different kind of embodiment is also found in studies investigating multimodal conversations, whereby gestures iconically support language functions (Louwerse *et al.* 2006). Eye gaze, facial expressions and hand gestures provide a wealth of examples of how language is often grounded in bodily experiences. Louwerse and Bangerter (2005) showed how language and gesture interact, and can even substitute each other. The use of these bodily experiences in communication facilitates thinking (Goldin-Meadow 2003; McNeill 1992). Gesture and speech are thereby co-expressive manifestations of one integrated system and form complementary components of one underlying process. In summary, arguing that language comprehension is not embodied would be analogous to arguing that language comprehension is not, to at least some extent, symbolic.

15.3 **Language comprehension is symbolic**

There is equally compelling evidence that symbols can derive meaning on the sole basis of other amodal symbols (Fodor 1975; Kintsch 1998; Landauer and Dumais 1997; Pylyshyn 1984). These symbolic approaches assume that conceptual representations are typically nonperceptual. That is, they are autonomous symbolic representations. Much of the recent evidence for the symbolic account of language processing comes from computer models. Two symbolic models of language processing have particularly been used in a applications and experiments. Models like HAL and LSA are computational models that determine the semantic relatedness by the frequency of word co-occurrences, as well as by the similarity of the contexts in which they co-occur (Landauer and Dumais 1997; Kintsch 1998a). Thus, these models compute higher-order relationships between text units such as words, sentences, paragraphs, and texts.

There is a large body of evidence showing that the output of these models correlate with human performance. For instance, HAL reliably categorized types of nouns (Lund and Burgess 1996), clustered common noun/proper name distinctions (Burgess and Conley 1998) as well as abstract or concrete words (Burgess and Lund 1997). LSA, another amodal symbolic system, does an impressive job on a variety of language tasks. For instance, by comparing student essays with ideal essays, LSA can automatically grade assignments (Landauer et al. 1998) and performs equally well as students on the Test of English as a Foreign Language (Landauer and Dumais 1997). LSA can even measure the coherence of texts (Foltz et al. 1998). Recently, LSA has been used as a measure of coherence in Coh-Metrix (Graesser *et al.* 2004; Louwerse *et al.* 2004), an internet-based tool that analyses texts on over 50 types of cohesion relations and over 200 measures of language, text, and readability.

LSA has also been used in intelligent tutoring systems like AutoTutor and iSTART. AutoTutor engages the student in a conversation on a particular topic like conceptual physics or computer literacy. AutoTutor uses LSA for its world knowledge and determines the semantic associations between a student answer, and ideal good and bad answers (Graesser *et al.* 2000; Graesser, Chapter 3, this volume). iSTART uses LSA in its teaching of reading strategies to students by providing appropriate feedback to students' self-explanations (McNamara *et al.* 2004). LSA is able to 'understand' metaphors (Kintsch 2000), and can assist in problem solving (Quesada *et al.* 2002). These studies suggest that meaning of language can be derived merely through the interrelations of words, sentences, and paragraphs.

15.4 **Language comprehension is both embodied and symbolic**

It seems obvious: the various examples shown so far suggest that neither an embodied nor a symbolic account alone can explain language comprehension, but that it is a convergence of the two accounts. However, if language comprehension is embodied and symbolic, how can this convergence be explained? Based on various theories, most notably Deacon (1997), Louwerse (2007) proposed a symbol interdependency hypothesis, making the argument that symbols can, but do not always have to, be grounded. That is, language is structured in such a way that many relationships that can also be found in the embodied world are structured in language. Language thereby provides a shortcut to the embodied relations in the world.

This claim is illustrated using Deacon's (1997) *language evolution theory* that proposes that icons, indices, and symbols are hierarchically structured. Deacon draws upon Peirce's (1923) theory of signs. A sign can be an icon, index, or symbol. An icon is mediated by a physical similarity between the sign and the object it refers to, as is the case in a picture representing an object. Icons require limited interpretation: we 're-cognize' what the sign stands for. Different from an icon, an index is mediated by spatial or temporal contiguity between sign and object, like smoke indicating the presence of fire. Note that there is no physical resemblance between smoke and fire, and that at least two iconic relationships are needed in order to build the indexical relationship. That is, the indexical link comes about through the iconic relationship between the representation of smoke and its referent, as well as the iconic relationship between the representation of fire and its referent. This indexical link between smoke and fire needs to be learnt at a certain time in a certain space and can be learnt through straightforward learning processes, as in classical and operational conditioning. Finally, a symbol is mediated by a conventional relationship between sign and object, like a yellow wristband as a symbol for hope, courage, and perseverance. Most of language belongs to this latter category. Whether we call an elephant an 'elephant' or a 'bulb' merely depends on what a language community agrees upon, not on a physical resemblance or a temporal and spatial contiguity.

Deacon now argues that icons, indices, and symbols represent levels of representation whereby indices are derived from icons (which correspond to objects in the real world) and symbols are derived from indices. The correspondence between words and objects thereby becomes a secondary relationship. But words do not only (indirectly) refer to objects, they are also linked to each other. Utilizing this network of symbolic relationships is what makes humans the only symbolic species. Whereas most species do not have the neurological architecture to process information beyond iconic and indexical levels, primates are able to recognize symbols, but only humans use the symbol system called language to its full extent.

The symbol interdependency hypothesis follows Deacon's taxonomy of icons, indices, and symbols in the sense that it recognizes that symbols can refer to each other, but are also connected as well to more concrete entities. This interdependency of symbols – to both objects and symbols – allows us to ground some words without the necessity to ground others, providing comprehenders with a convenient mnemonic support for language processing. Comprehenders do not have to activate every single detail of the embodied representations, but they can use a simple shorthand. Using a neural network metaphor we would say that all nodes related to these symbols – including embodied representations – receive some activation. However, by virtue of being located far away from the symbols as implied by Deacon (1997)'s theory – separated by several intermediate layers of neurons – the embodied representations by default don't get enough activation to be in working memory. That is, in a default conversation and reading setting, whereby comprehenders quickly process language, they do not have to perceptually simulate each and every word they encounter. That some particular words trigger visual/motor imagery or other embodied representations is common in certain circumstances, but it is argued that efficiency in language processing simply does not yield a continuous systematic grounding.

Similar to other accounts, the symbol interdependency hypothesis implies the transduction of information into a new representation. Take for instance Barsalou's (1999) theory of *perceptual symbols*. According to Barsalou, the brain creates neural records during perception. Subsets of these perceptual states are extracted and stored in long-term memory. These subsets of perceptual states are what Barsalou calls perceptual symbols. Thus, knowledge is not transduced into an amodal symbolic representation, but instead into modality-specific states (e.g., visual states). However, during retrieval these perceptual states can function symbolically. Barsalou's point is that a physical stimulus (e.g., seeing a car) is transduced into modality-specific states (e.g., sensory representation of a car), rather than an amodal representation outside the sensorimotor system. It seems more convenient to remember a smell as a smell (with associative sensory representations) and a touch as a touch. Similarly, the symbol interdependency hypothesis states that linguistic expressions by default get transduced into a symbolic representation, because it is generally more efficient to do so. It seems more convenient to remember an amodal symbol as an amodal symbol (with associative symbolic representations).

15.5 **Activation of embodied representations**

According to the symbol interdependency hypothesis, comprehenders can bootstrap the meaning of amodal symbols through their interrelations and don't need to ground all amodal symbols under all circumstances. There may be theoretical evidence for this claim, but the question needs to be answered how this claim can explain the wealth of evidence for embodied representations. Dozens of psycholinguistic studies have been reported over the last half decade – some mentioned above – all finding evidence for the activation of embodied representations in language comprehension. How can this wealth of evidence warrant the claim that embodied representations are not always sufficiently activated under all circumstances? We have two answers to this question.

First, in some psycholinguistic embodiment experiments participants are typically cued to activate embodied representations like images or motor skills. These experiments demonstrate quite powerfully that cognition cannot be strictly symbolic, because non-linguistic modalities get activated. However, these experiments do not show that no amodal representations get activated. For instance, Stanfield and Zwaan (2001) asked participants to verify whether a picture depicted an object in a sentence. Participants responded more quickly to a picture of an eagle with its wings spread out after reading 'The ranger saw an eagle in the sky' than after 'The ranger saw an eagle in the tree' (Zwaan *et al.* 2002) or to a picture of a horizontally depicted nail after reading 'He pounded the nail into the wall' than to a vertically depicted nail, but *vice versa* after reading 'He pounded the nail into the floor' (Stanfield and Zwaan 2001). According to the authors, understanding the sentence activates pictorial information (i.e., an embodied representation) matching the content of that sentence. Consequently, the reaction times are shorter when the picture matches this pictorial information. Similarly, according to Glenberg and Kaschak's (2002) ACE paradigm there is a compatibility effect between the direction described linguistically in a sentence and the action executed motorically. When subjects are presented with a sentence like 'Courtney handed you the notebook' subjects are faster in making sensibility judgements when the response button is close to them, but when a sentence like 'You handed Courtney the notebook' is presented subjects are faster in making the judgement when the response button is away from them. Both Stanfield and Zwaan (2001) and Glenberg and Kaschak (2002) show that it is more efficient to transduce a physical stimulus in a modality-specific state rather than transducing it into an amodal symbolic representation (Barsalou 1999). That is, these studies show that embodied representations are activated when comprehenders are cued to activate them, but not that amodal symbolic representations do not get activated.

Secondly, in some psycholinguistic embodiment experiments, the experimental instructions activate deep levels of processing, whereby we consider depth in terms of the extent to which meaningfulness is extracted from the stimulus (Lockhart and Craik 1990). Interestingly, as we will show later, most evidence for activation of embodied representations comes from experiments whereby participants are involved in tasks that require cognitive processing deeper than what would typically be required in default language comprehension. For instance, Spivey and Geng (2001) found that subjects acted out the mental image of a passage they read. Subjects listened to a story that had descriptions of upward, downward, leftward, and rightward events, like in the following text:

> Imagine that you are standing across the street from a 40-storey apartment building. At the bottom there is a doorman in blue. On the 10th floor, a woman is hanging her laundry out of the window. On the 29th floor, two kids are sitting on the fire escape, smoking cigarettes. On the very top floor, two people are screaming.

Unbeknownst to participants, their eye movements were recorded as they followed the text. Spivey and Geng found that eye movements were in the direction of the described directions (vertical in this case), suggesting that lower-level motor processes are activated with higher-level cognitive processes. The story presented to participants required considerable cognitive effort. The question could therefore be raised whether a text not cueing the participant to think deeply through the materials, as in the average narrative text, may have led to the same mental imagery. Though this question remains unanswered for now, it is noteworthy that many embodiment experiments require the participant to deeply process the materials, more than what would be needed to assign meaning to language, as we will argue next.

In summary, embodiment experiments clearly show that language comprehension cannot be solely symbolic. However, they do not show that language comprehension is solely embodied. Comprehenders are often explicitly cued to activate nonlinguistic modalities or are involved in a deep semantic analysis task whose nature is such that different kinds of representations – modal and amodal – are activated. Roughly stated, psycholinguistic embodiment experiments can be classified by task and nature of the stimuli. Table 15.1 gives an overview of prominent psycholinguistic embodiment experiments that provide evidence for embodied representations and the associated experimental tasks. Though this may not be a complete overview, we believe it is an adequate representation of psycholinguistic studies in embodied cognition.

The table indicates that the number of embodiment experiments whereby deeper semantic analysis is needed (semantic judgement, sensibility ratings, memory load) or whereby nonlinguistic stimuli are used or motor skills are required (pictorial information, movement to button press) outweigh the studies whereby solely linguistic information is used and shallow language processing suffices. The purpose of this table is not to be exhaustive, to debate whether a study uniquely belongs in a cell, or even whether studies can be classified solely on the basis of these two dimensions. Instead, the table poses the question under what circumstances embodied representations get activated and whether embodied representations supersede symbolic representations in standard (shallow linguistic) language tasks.

Table 15.1 Classification embodiment studies

	Linguistic stimuli only	**Linguistic and nonlinguistic stimuli**
'Shallow' processing	Pecher *et al*. (2003); 1,2	Fincher-Kiefer (2001); 2 Matlock and Richardson (2004) Richardson *et al*. (2001); 1 Richardson *et al*. (2003); 1,2 Stanfield and Zwaan (2001) Zwaan *et al*. (2004) Zwaan *et al*. (2002); 1,2 Zwaan and Yaxley (2003); 1–3
'Deep' processing	Borghi and Barsalou (2001) Borghi (2004); 1–5 Borghi *et al*. (2004a); 1,2 Demarais and Cohen (1998); 1 Fincher-Kiefer (2001); 1 Glenberg and Robertson (2000); 1–3 Matlock *et al*. (2005); 1–3 Richardson *et al*. (2001); 2 Spivey *et al*. (2000) Spivey and Geng (2001); 1 Wu and Barsalou (2001)	Bergen and Wheeler (2005); 1, 2 Demarais and Cohen (1998); 2 Borghi *et al*. (2004a); 3 Boroditsky (2000); 1–3 Boroditsky and Ramscar (2002); 1–4 Glenberg and Kaschak (2002); 1,2AB Kaschak *et al*. (2005); 1,2 Spivey and Geng (2001); 2 Tseng and Bergen (2005)

N.B. The number following the author and year refers to the experiment in the study.

15.5.1 Activating embodied relations under certain conditions

What is the evidence that in many circumstances embodiment representations are not necessarily activated? Is there evidence that comprehenders can actually rely on the organization of the symbolic system to derive meaning? The symbol interdependency hypothesis makes two predictions. First, because symbols are interdependent on one another, it seems convenient for language users to access language symbolically, unless embodied representations are cued by the task, visual imagery, or motor movements. Second, because language develops through our interaction with the physical world, and because of its efficiency embodied relations are encoded in language structures. In the remainder of this paper we will give an overview of the evidence for these two predictions.

There are various experiments whereby the participant is exposed to linguistic information only, and whereby the task does not explicitly trigger mental imagery. One such study by Zwaan and Yaxley (2003) utilized a spatial Stroop test in which participants saw word-pairs that were either presented on the vertical axis in their iconic order ('attic' above 'basement') or in a reverse-iconic order ('basement' above 'attic'). Zwaan and Yaxley found that participants were faster in making a semantic judgement for the iconic than for the reverse-iconic pairs. They argued that this proves that embodied representations are activated during language processing. According to the symbol interdependency hypothesis, embodied relations are activated under certain conditions, like the 'deeper' processing of the stimuli or a processing such that the activation of embodied representations is desirable.

Intrigued by the interpretation of Zwaan and Yaxley's (2003) findings, Louwerse and Jeuniaux (under review) repeated the experiment by varying a series of factors, one of which was the semantic relation between the two items. Motivated by the idea that something intrinsically symbolic would also have a role to play, we split all pairs in function of the degree of semantic relationship of their items as determined by LSA. This semantic separation was realized by computing the LSA cosine value for each pair. We had therefore items with a strong semantic relation ('airplane'–'runway') and items with a weak semantic relation ('penthouse'–'lobby'). Besides these experimental pairs, we added filler pairs composed of one or two nonwords. Secondly, we varied the task. In the first experiment we used the same semantic judgement task implemented by Zwaan and Yaxley. In the second experiment we used a lexical decision task, where participants decided whether the pair consisted of words only or included nonwords. The idea behind varying the task was that in both cases some semantic processing would be involved but with different degrees of intensity ('deeper' processing in the case of semantic judgement and 'shallower' processing in the case of lexical judgement), and that the activation of embodied representations would be more likely in the former than in the latter, thereby illustrating that these representations are not activated under all circumstances involving semantic processing.

What we found corresponds to this line of reasoning: an effect for iconicity was found in the semantic judgement task but not in the lexical decision task. In the semantic task, as in Zwaan and Yaxley (2003)'s experiment, the participants were faster at judging the semantic similarity of two items if they were located in iconic order on the vertical axis. We concluded that for deeper levels of processing – like semantic judgement, where processing times tend to be higher than for shallower levels of processing like lexical decision – embodied relations would be likely more activated. Of course, the argument could be made that in a shallow language processing task like lexical decision where participants merely decide whether a word is a word or a non-word no real language processing takes place. Interestingly, however, this argument does not hold since we found semantic effects in both experiments (see also Meyer and Schvaneveldt 1971). Indeed, the judgement of semantic relatedness and the response times for lexicality were shorter when the two items were strongly related semantically. Therefore, comprehension *must* have taken place, but at a level shallower (and faster) than in the semantic judgement task.

In a follow-up experiment, we presented the word-pairs horizontally, as did Zwaan and Yaxley (2003). For-word pairs presented horizontally, semanticity but not iconicity effects were found. One explanation is that word-pairs presented horizontally are not perceptually simulated in the same way as word-pairs presented horizontally. An alternative explanation that we pursued is that vertically presented items seem to require the comprehender to process the words independently, whereas with horizontally presented items the words are integrated into one representation (Bradshaw *et al.* 1981). These data do not support the claim that no embodied representations are formed. Instead, they favour the idea that both symbolic and embodied representation are at the core of language processing. The data also suggest that the consideration of embodied relationships slows down the processing of language and therefore may not always occur by default in language processing.

15.6 **Language encodes relations in the world**

If relations between amodal symbols provide a convenient shortcut to understanding the world described by the language, the structure of the amodal symbol system would be most efficient if it somehow encodes the relations in the physical world. Louwerse (2007) argues that it is indeed likely that language has been shaped by thought. For instance, phrasal structures match perceptual concepts, with adjectives being placed right before or after a noun, so that the 'grey mouse' is easily understood as a mouse that is grey. This even applies to languages that use case suffixes like German, Latin, and Turkish, and would allow for adjectives to be detached from their nouns. Similarly, when we think of actions, we think of subjects (often agents) and objects (often patients), most likely in that order. Indeed, in the vast majority of the world's languages, subjects precede objects (Greenberg 1963).

In a number of other studies we have further explored the extent to which language structures encode perceptual information. In some we used n-grams to compute the frequency with which words occur in a certain order; for others we had to rely on higher-order relations between words using LSA to avoid problems of data sparsity. LSA estimates semantic similarities on a scale of –1 to 1 between the latent semantic structure of words and linguistic contexts. The input to LSA is a set of corpora segmented into contexts like sentences, paragraphs, or whole texts. Mathematical transformations create a large multidimensional words–contexts matrix from the input. By removing dimensions corresponding to small variance and keeping the dimensions corresponding to larger variance, the representation of each term is further reduced as a smaller vector with only *k* dimensions. Each word now becomes a weighted vector on 300 or so dimensions. The semantic relationship between words can be estimated by taking the dot product (cosine) between two vectors.

Various studies have shown that LSA is able to accurately predict the semantic relations between words. Louwerse *et al.* (2006) have shown that LSA is also able to predict temporal and spatial relations. For instance, LSA in combination with multidimensional scaling (MDS) algorithms was able to accurately order the days of the week as well as the months of the year. Moreover, it was able to determine the order of time distances of relative time units ('second', 'minute', 'hour', 'day', 'week', 'month', 'year') and pointing time adverbs ('ago' and 'later'). LSA was also able to capture world knowledge by predicting the distances between spatial coordinates. Cosine values were higher when comparing 'New York' and 'Boston' than for 'New York' and 'Seattle' because, in natural discourse, place names that are closer in proximity occur more frequently in the same documents and, similarly, their contexts are shared across multiple documents. This finding emerged in the comparison of cosines and physical distances of 28 of the largest cities in the United States, the largest cities in the world, and even for different languages (French and English).

We have also conducted a number of studies to determine to what extent symbolic algorithms can explain data found in embodiment experiments. Louwerse (2007) for instance used the same technique as Louwerse *et al.* (2006) computing cosine values from the Glenberg and Robertson (2000) data and analyzing the results using MDS. Glenberg and Robertson presented subjects with a setting sentence ('After wading barefoot in the lake, Erik needed something to get dry'), followed by a related sentence ('He used his towel to dry his feet'), an afforded sentence ('He used his shirt to dry his feet) or a nonafforded sentence ('He used his glasses to dry his feet'). Sensibility and envisioning data from human subjects showed a difference between related/afforded versus nonafforded sentences with lowest scores for the latter. Glenberg and Robertson showed that no such difference was found by comparing the LSA cosine values. However, Louwerse (2007) showed that the combination of LSA with MDS showed no significant differences between the related and the afforded sentences, and an expected difference between the related sentences and the nonafforded sentences, with the related sentences yielding higher scores than nonafforded sentences. Only a marginally significant difference was found between afforded and nonafforded sentences. A correlation between the nonafforded sentence and the setting ('dry'–'glasses'), however, showed a strong negative correlation, suggesting that these are furthest away in terms of Euclidean distance. Furthermore, when the MDS coordinates were compared with Glenberg and Robertson's sensibility and envisioning ratings, a positive correlation was found. These findings show that the embodiment differences could also be obtained using symbolic approaches, suggesting an encoding of embodied relations in language structures.

Similarly, Louwerse (under review) reanalysed data from Borghi *et al.* (2004) and Zwaan and Yaxley (2003) using corpus linguistic methods. For instance, Borghi *et al.* tested concept sets that consisted of an orientation sentence ('There is a doll standing on a table in front of you') followed by a probe word. Two probe words were related to the concept noun of the sentence ('doll'), one describing the upper part of the concept ('hair') the other describing the lower part ('ankle'). Participants were asked to verify whether the probe word was part of the concept by pressing 'yes' or 'no'. The beauty of the experiment lay in the position of the 'yes' and 'no' buttons. In half of the experiments 'yes' was located above 'no' (yes-is-up condition); in the other half the position was reversed (yes-is-down condition). Participants in a movement condition reached out for the buttons, whereas those in a no-movement condition had their hands on the buttons already. Borghi *et al.*'s results showed that response times for the upper parts (e.g., 'hair') were shorter than for the lower parts (e.g., 'ankle') in the yes-is-up condition in the movement condition only. We replicated Borghi *et al.*'s movement results in a no-movement condition by presenting participants with vertically presented word-pairs ('doll'–'hair' and 'doll–ankle'). Incidentally, in a corpus linguistic analysis we found that the frequency of concept–upper-part word-pairs is higher than that of concept–lower-part word-pairs, suggesting that upper–lower part distinctions found in the physical world can also be found in language structures. A similar pattern was found when an experiment replicated Zwaan and Yaxley's findings and word-order frequencies were computed for these stimuli. Iconic word order ('attic' preceding 'basement') was more frequent than reverse-iconic word order ('basement' preceding 'attic'). Moreover, word-order frequency explained more of the variance of the reaction time data than did iconicity. These findings do not falsify the results of Borghi *et al.* or Zwaan and Yaxley, but they do suggest that the iconic representations of concepts in the world are also encoded in language structures. The notion of thought shaping language should at least be considered when debating the extent to which language comprehension is embodied and symbolic.

15.7 **Conclusion**

There is evidence that language comprehension is embodied. There is also evidence that language comprehension is symbolic. The symbol interdependency hypothesis argues that it is both: symbols link to other symbols through higher-order relationships and they also refer to objects in the physical world. In many language tasks, comprehenders can draw meaning from symbolic relationships without any extensive grounding being needed. In different circumstances, due to the nature of the instructions or the stimuli; embodied representations become more desirable or useful. The impression may have been given that our account is polemic in nature. It is not. Instead, it argues that debate whether language comprehension is symbolic or embodied should make room for more timely research questions. These include under what circumstances embodied and symbolic representations are activated, and to what extent language shapes thought and thought shapes language.

Debate

Julio Santiago: A really interesting talk, thank you. What I don't really get is the last part in which you compared the horizontal position of word-pairs and argue that the horizontal presentation yields easier processing than vertical presentation. Zwaan and Yaxley would argue that this is exactly the evidence for perceptual simulations.

Max Louwerse: But my point is that when you present items horizontally, you will find shorter reaction times in the first place compared with when you present items vertically. So, there may be different psycholinguistic processes if you process words vertically or horizontally. For instance, it may be that if you process items horizontally, the activation of one word immediately affects the activation of the second word, whereas in vertical processing this carry-over effect is less salient. That is, in vertical processes, the items in the word pairs may be processed more independently. That is at least a possible explanation.

Santiago: But it still doesn't give an explanation for the effect that you find when the words are presented in the vertical dimension.

Louwerse: Well, actually, we ran a multiple regression analysis on the reaction-time data we got from the semantic judgement task and lexical decision experiments, dummy coding for word order. That is, we had two variables in the regression analysis, one (embodied) iconicity and one (symbolic) word order. And what we found was that only 2% of the data in both cases can be explained, but the data are better explained by word order than by iconicity.

Robert Goldstone: I'm interested in the experiments that you talked about at the beginning, where you showed that you could reconstruct the spatial configuration of the days of the week, months of the year, or US cities. I mean, I guess, part of my skepticism is wondering whether it's really just a sort of a side effect, That there's something going on in people's minds that would make them talk about Los Angeles at the same time they would talk about San Francisco because of their spatial proximity.

Louwerse: Exactly.

Goldstone: So it makes sense that you could then do some statistical analysis to reconstruct the underlying space that was in the mind of the producer when they were creating it, but I guess I'm still holding out for wanting to have a sort of communal process that explains both the production and the comprehension processes, in order to facilitate communication between them. So, I guess my question to you is, do you see the same thing going on at the production stage? In other words, is there no spatial grounding when they are talking about these cities together?

Louwerse: I'm not arguing against spatial grounding, but, according to what I call the symbol independency hypothesis, you don't have to ground every single item. You only have to ground one or two, and that will bootstrap the grounding of all the others. That's basically what I argue for. So, with regards to what you say about cities, yes we talk about Los Angeles more in the context of San Francisco than New York, for instance. I think that's the whole point. Apparently language has those structures encoded based on the fact that the physical world is structured this way.

Walter Kintsch: So, the series is incomplete. It is a series of comprehension: it does not explain production. And I guess it's embodied. The reason why San Francisco and Los Angeles are closer together in LSA is because people live in the world where they are more likely to talk about those two things. So you can't understand new things by making new combinations of all things. But it doesn't explain production.

Arthur Glenberg: Marvellously creative, thank you. Of course I've got a lot of questions about details but we can do that later. The main question I want to pose is: LSA depends on what's already been done, what's already been written. And, it then seems, if we're going to depend completely on LSA we can never be creative, we can never use language to describe something new. So, do you think that whenever we are using language to describe a new experience we must be embodied?

Louwerse: Again, I'm not arguing against embodiment. With regards to your question about novel situations, I don't know. But LSA could still apply to novel situations. For instance, we could run an experiment where we replaced some existing words with nonwords – that would be a sort of novel scenario I guess – and determine the meaning of the new words. If you replace existing words with nonwords, the nonwords would acquire semantic information from the other real words. I think that sort of relates to the novel situation.

Glenberg: I meant something more along the lines of, well, not learning a novel word but, for instance, telling somebody about my experience in Garachico. And you know, assuming that Garachico is represented in a TASA corpus, would I then be able to describe to you the layout of the town so that somebody could then know how to get around the town?

Louwerse: First of all, when you're talking about a new place you have to at least know the place; information about that place should be somewhere in your corpus. I would then argue that if there is enough contextual information, you would indeed be able to draw a map. In fact, that's what we have done in Louwerse *et al.* (2006) and what we are doing in current studies. Granted, it's true that LSA usually works backwards and needs lots of context. But when it does have that context I don't really see why it would not be able to make predictions for the future, for novel situations.

Author note

This research was supported by National Science Foundation grant NSF-IIS-0416128. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding institution. Correspondence concerning this article should be addressed to Max M Louwerse, Institute for Intelligent Systems / Department of Psychology, University of Memphis, 202 Psychology Building, Memphis, TN 38152-3230, USA.

References

Altmann GTM, Kamide Y (1999). Incremental interpretation at verbs: restricting the domain of subsequent reference. *Cognition*, 73, 247–64.

Barsalou LW (1999). Perceptual symbol systems. *Behavior and Brain Sciences*, 22, 577–660.

Bergen B, Wheeler K (2005). Sentence understanding engages motor processes. In *Proceedings of the 27th Annual Conference of the Cognitive Science Society*. Mahwah, NJ: Erlbaum.

Bishop M, Preston J (2001). *Essays on Searle's Chinese Room Argument*. Oxford: Oxford University Press.

Borghi AM (2004). Objects concepts and action: extracting affordances from objects' parts. *Acta Psychologica*, 115, 69–96.

Borghi AM, Barsalou LW (2001). Perspectives in the conceptualization of categories. *Meeting of the Psychonomic Society*, Orlando, FL, November 2001.

Borghi AM, Caramelli N, Setti A (2004). Conceptual information on objects' location. *Brain and Language*, 93, 140–51

Borghi AM, Glenberg AM, Kaschak MP (2004a). Putting words in perspective. *Memory and Cognition*, 32, 863–73.

Boroditsky L (2000). Metaphoric structuring: understanding time through spatial metaphors. *Cognition*, 75, 1–28.

Boroditsky L and Ramscar M (2002). The roles of body and mind in abstract thought. *Psychological Science*, *13*, 185-188.

Bradshaw JL, Nettleton NC, Taylor MJ (1981). The use of laterally presented words in research intocerebral asymmetry: is directional scanning likely to be a source of artifact? *Brain and Language*, 14, 1–14.

Bransford JD, Franks JJ (1971). The abstraction of linguistic ideas. *Cognitive Psychology*, 2, 331–50.

Burgess C (1998). From simple associations to the building blocks of language: Modeling meaning in memory with the HAL model. *Behavior Research Methods, Instruments, and Computers*, 30, 188–98.

Burgess C, Conley P (1998). Developing semantic representations for proper names. *Proceedings of the Cognitive Science Society* (pp. 185–90). Hillsdale, NJ: Erlbaum.

Burgess C, Lund K (2000). The dynamics of meaning in memory. In E Dietrich, A Markman, Eds. *Cognitive Dynamics: Conceptual Change in Humans* and *Machines* (pp. 295–342). Mahwah, NJ: Erlbaum.

Burgess C, Lund K (1997). Modeling parsing constraints with high-dimensional context space. *Language and Cognitive Processes*, 12, 177–210.

Cangelosi A, Harnad S (2000). The adaptive advantage of symbolic theft over sensorimotor toil: grounding language in perceptual categories. *Evolution of Communication*, 4, 117–42.

Chambers CG, Tanenhaus MK, Eberhard KM, Filip H, Carlson GN (2002). Circumscribing referential domains in real-time language comprehension. *Journal of Memory and Language*, 47, 30–49.

Deacon T (1997). *The Symbolic Species: The Co-evolution of Language and the Human Brain*. London: Allen Lane.

Demarais A, Cohen BH (1998). Evidence for image-scanning eye movements during transitive inference. *Biological Psychology*, 49, 229–47.

Fincher-Kiefer R (2001). Perceptual components of situation models. *Memory and Cognition*, 29, 336–43.

Fodor JA (1975). *The Language of Thought*. Cambridge, MA: Harvard University Press.

Foltz PW, Kintsch W, Landauer TK (1998). The measurement of textual coherence with latent semantic analysis. *Discourse Processes*, 25, 285–307.

Glenberg AM, Kaschak MP (2002). Grounding language in action. *Psychonomic Bulletin and Review*, 9, 558–65.

Glenberg AM, Robertson DA (2000). Symbol grounding and meaning: a comparison of high-dimensional and embodied theories of meaning. *Journal of Memory and Language*, 43, 379–401.

Glenberg AM (1997). What memory is for. *Behavioral and Brain Sciences*, 20, 1–55.

Goldin-Meadow S (2003). *Hearing gesture: How our* Hands *Help us Think*. Cambridge, MA: Harvard University Press.

Graesser AC, McNamara DS, Louwerse MM, Cai Z (2004). Coh-Metrix: analysis of text on cohesion and language. *Behavior Research Methods, Instruments, and Computers*, 36, 193–202.

Graesser AC, Wiemer-Hastings P, Wiemer-Hastings K, Harter D, Person N; Tutoring Research Group (2000). Using latent semantic analysis to evaluate the contributions of students in AutoTutor. *Interactive Learning Environments*, 8, 128–48.

Greenberg J, Ed (1963). *Universals of Language*. Cambridge, MA: MIT Press.

Harnad S (1990). The symbol grounding problem. *Physica D*, 42, 335–46.

Howard MW, Kahanna MJ (2002). When does semantic similarity help episodic retrieval? *Journal of Memory and Language*, 46, 85–98.

Kaschak MP, Madden CJ, Therrriault DJ *et al*. (2005). Perception of motion affects language processing. *Cognition*, 94, B79–89.

Kintsch W (2000). Metaphor comprehension: a computational theory. *Psychonomic Bulletin* and *Review*, *7*, 257-266.

Kintsch W (1998a). *Comprehension: A Paradigm for Cognition*. New York, NY: Cambridge University Press.

Kintsch W (1998b). The representation of knowledge in minds and machines. *International Journal of Psychology*, 33, 411–20.

Kintsch W, Keenan J (1973). Reading rate and retention as a function of the number of propositions in the base structure of the text. *Cognitive Psychology*, 5, 257–74.

Klatzky RL, Pellegrino JW, McCloskey DP, Doherty S (1989). Can you squeeze a tomato? The role of motor representations in semantic sensibility judgements. *Journal of Memory and Language*, 28, 56–77

Landauer TK, Dumais ST (1997). A solution to Plato's problem: the latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, 104, 211–40.

Landauer TK, Foltz PW, Laham D (1998a). An introduction to latent semantic analysis. *Discourse Processes*, 25, 259–84.

Landauer TK, Laham D, Foltz PW (1998b). Learning human-like knowledge by singular value decomposition: a progress report. In MI Jordan, MJ Kearns, SA Solla, Eds. *Advances in Neural Information Processing Systems 10* (pp. 45–51). Cambridge, MA: MIT Press.

Landauer TK, McNamara DS, Dennis S, Kintsch W, Eds. (2007). *Handbook of Latent Semantic Analysis*. Mahwah, NJ: Erlbaum.

Landauer TK (2002). On the computational basis of learning and cognition: arguments from LSA. In N Ross, Ed. *The Psychology of Learning and Motivation*, 41, 43–84.

Louwerse MM (2007). Iconicity in amodal symbolic representations. In T Landauer, D McNamara, S Dennis, W Kintsch, Eds. *Handbook of Latent Semantic Analysis*. Mahwah, NJ: Erlbaum.

Louwerse MM (under review). Min(d)ing symbols. Embodied representations are encoded in language.

Louwerse MM, Bangerter A (2005). Focusing attention with deictic gestures and linguistic expressions. In B Bara, L Barsalou, M Bucciarelli, Eds. *Proceedings of the Cognitive Science Society* (pp. 1331–6). Mahwah, NJ: Erlbaum.

Louwerse MM, Cai Z, Hu X, Ventura M, Jeuniaux P (2006). Cognitively inspired natural-language based knowledge representations: further explorations of latent semantic analysis. *International Journal of Artificial Intelligence Tools*, 15, 1021–39.

Louwerse MM, Jeuniaux P (under review). Symbolic language in an embodied world.

Louwerse MM, Jeuniaux P, Hoque ME, Wu J, Lewis G (2006). Multimodal communication in computer-mediated map task scenarios. *Proceedings of the 27th Annual Conference of the Cognitive Science Society*. Mahwah, NJ: Erlbaum.

Louwerse MM, McCarthy PM, McNamara DS, Graesser AC (2004). Variation in language and cohesion across written and spoken registers. In K Forbus, D Gentner, T Regier, Eds. *Proceedings of the 26th annual conference of the Cognitive Science Society* (pp. 843–8). Mahwah, NJ: Erlbaum.

Lund K, Burgess C (1996). Producing high-dimensional semantic spaces from lexical co-occurrence. *Behavior Research Methods, Instrumentation, and Computers*, 28, 203–8.

Matlock T, Richardson DC (2004). Do eye movements go with fictive motion? In K Forbus, D Gentner, T Regier, Eds. *Proceedings of the 26th Annual Conference of the Cognitive Science Society*. Mahwah, NJ: Erlbaum.

Matlock T, Ramscar M, Boroditsky L (2005). The experiential link between spatial and temporal language. *Cognitive Science*, 29, 655–64.

McNamara DS, Levinstein IB, Boonthum C (2004). iSTART: Interactive strategy trainer for active reading and thinking. *Behavioral Research Methods, Instruments, and Computers*, 36, 222–33.

McNeill D (1992). *Hand and Mind: What Gestures Reveal about Thought*. Chicago, IL: University of Chicago Press.

Meyer DE, Schvaneveldt RW (1971). Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. *Journal of Experimental Psychology*, 90, 227–34.

Ogden CK, Richards IA (1923). *The Meaning of Meaning*. London: Routledge & Kegan Paul.

Pecher D, Zwaan RA, Eds (2005). *Grounding Cognition: The Role of Perception and Action in Memory, Language, and Thinking*. New York, NY: Cambridge University Press.

Pecher D, Zeelenberg R, Barsalou LW (2003). Verifying conceptual properties in different modalities produces switching costs. *Psychological Science*, 14, 119–24.

Peirce CS (1931–58). In C Hartshorne, P Weiss, A Burks, Eds. *The Collected Papers of Charles Sanders Peirce*. Cambridge, MA: Harvard University Press.

Plato (1902). *The Dialogues of Plato. Translated into English with Analyses* and *Introductions* (BMA Jowett, Trans.). New York, NY: Charles Scribner's Sons. (Original work published 360BC).

Pulvermüller F (1999). Words in the brain's language. *Behavioral and Brain Sciences*, 22, 253–70.

Pylyshyn ZW (1984). *Computation and cognition: towards a foundation for cognitive science*. Cambridge, MA: MIT Press.

Quesada JF, Kintsch W, Gomez E (2003). Latent problem solving analysis as an explanation of expertise effects in a complex, dynamic task. In R Alterman, D Kirsh, Eds. *Proceedings of the 25th Annual Conference of the Cognitive Science Society*. Boston, MA.

Richardson DC, Spivey MJ, Cheung J (2001). Motor representations in memory and mental models: embodiment in cognition. *Proceedings of the 23rd Annual Meeting of the Cognitive Science Society* (pp. 867–72). Mawhah, NJ: Erlbaum.

Richardson DC, Spivey MJ, Barsalou LW, McRae K (2003). Spatial representations active during real-time comprehension of verbs. *Cognitive Science*, 27, 767–80.

Rosch E (1973). On the internal structure of perceptual and semantic categories. In TE Moore, Ed. *Cognitive Development* and *the Acquisition of Languag*e (pp. 111–44). New York, NY: Academic Press.

Roy D, Pentland A (2002). Learning words from sights and sounds: A computational model. *Cognitive Science*, 26, 113–46.

Searle JR (1980). Minds, brains, and programs. *Behavioral and Brain Sciences*, 3, 417–57.

Spence C, Nicholls MER, Driver J (2000). The cost of expecting events in the wrong sensory modality. *Perception and Psychophysics*, 63, 330–6.

Spivey M (2006). *The Continuity of Mind*. New York, NY: Oxford University Press.

Spivey MJ, Geng JJ (2001). Oculomotor mechanisms activated by imagery and memory: wye movements to absent objects. *Psychological Research*, 65, 235–41.

Spivey M, Tanenhaus M, Eberhard K, Sedivy J (2002). Eye movements and spoken language comprehension: effects of visual context on syntactic ambiguity resolution. *Cognitive Psychology*, 45, 447–81.

Spivey M, Tyler M, Richardson D, Young E (2000). Eye movements during comprehension of spoken scene descriptions. In *Proceedings of the 22nd Annual Conference of the Cognitive Science Society* (pp. 487–92). Mahwah, NJ: Erlbaum.

Spivey-Knowlton M, Tanenhaus M, Eberhard K, Sedivy J (1998). Integration of visuospatial and linguistic information in real-time and real-space. In P Olivier, K-P Gapp, Eds. *Representation and Processing of Spatial Expressions* (pp. 201–14). Mahwah, NJ: Erlbaum.

Stanfield RA, Zwaan RA (2001). The effect of implied orientation derived from verbal context on picture recognition. *Psychological Science*, 12, 153–6.

Steyvers M, Shiffrin RM, Nelson DL (2004). Word association spaces for predicting semantic similarity effects in episodic memory. In A Healy, Ed. *Experimental Cognitive Psychology and its Applications* (pp. 237–49). Washington, DC: American Psychological Association.

Tanenhaus MK, Spivey-Knowlton MJ, Eberhard KM, Sedivy JC (1995). Integration of visual and linguistic information in spoken language comprehension. *Science*, 268, 1632–4.

Tseng M, Bergen B (2005). Lexical processing drives motor simulation. In B Bara, L Barsalou, M Bucciarelli, Eds. *Proceedings of the Cognitive Science Society*. Mahwah, NJ: Erlbaum.

van Dijk TA, Kintsch W (1983). *Strategies of Discourse Comprehension*. New York, NY: Academic Press.

Vigliocco G, Vinson DP, Lewis W, Garrett MF (2004). Representing the meanings of object and action words: The featural and unitary semantic space hypothesis. *Cognitive Psychology*, *48*, 422-488.

Westermann G, Reck Miranda E (2004). A new model of sensorimotor coupling in the development of speech. *Brain and Language*, 89, 393–400.

Wiemer-Hastings K, Xu X (2005). Content differences for abstract and concrete concepts. *Cognitive Scienc*e, 29, 719–36.

Wu LL, Barsalou LW (under review). Grounding concepts in perceptual simulation. I: Evidence from property generation.

Wu LL, Barsalou LW (2004). Perceptual simulation in property verification. *Memory* and *Cognition*, 32, 244–59.

Zwaan RA (2004). The immersed experience: toward an embodied theory of language comprehension. In BH Ross, Ed. *The Psychology of Language* and *Motivation*, *44*. New York, NY: Academic Press.

Zwaan RA, Madden CJ, Yaxley RH, Aveyard ME (2004). Moving words: dynamic mental representations in language comprehension. *Cognitive Science*, 28, 611–19.

Zwaan RA, Stanfield RA, Yaxley RH (2002). Do language comprehenders routinely represent the shapes of objects? *Psychological Science*, 13, 168–71.

Zwaan RA, Yaxley RH (2003). Spatial iconicity affects semantic-relatedness judgements. *Psychonomic Bulletin* and *Review*, 10, 954–8.