

Putting words in perspective

ANNA M. BORGHI

University of Bologna, Bologna, Italy

and

ARTHUR M. GLENBERG and MICHAEL P. KASCHAK

University of Wisconsin, Madison, Wisconsin

In this article, we explore the nature of the conceptual knowledge retrieved when people use words to think about objects. If conceptual knowledge is used to simulate and guide action in the world, then how one can interact with an object should be reflected in the speed of retrieval and the content that is retrieved. This prediction was tested in three experiments in which a part verification procedure was used. Experiments 1 and 2 demonstrated that speed of part verification varied with the perspective imposed on the object by the language used to name the object (e.g., “You are driving a car” or “You are fueling a car”). In Experiment 3, parts were chosen so that actions directed toward them (on the real object) require movement upward (e.g., the roof of a car) or downward (e.g., the wheels of a car). Orthogonally, responding “yes” required an upward or a downward movement to a response button. Responding in a direction incompatible with the part’s location (e.g., responding downward to verify that a car has a roof) was slow relative to responding in a direction compatible with the part’s location. These results provide a strong link between concept knowledge and situated action.

The knowledge retrieved when object names are used to guide thinking about the relevant objects depends on the context in which the names are presented (see, e.g., Anderson & Ortony, 1975; Roth & Shoben, 1983). The experiments presented in this article are aimed at testing a specific explanation for this phenomenon: that conceptual knowledge is rooted in perception- and action-based representations, which are used to simulate and guide action in the world. This view suggests that object parts are often retrieved in a manner reflecting their spatial organization, because spatial organization plays a central role in how one can interact with the object in the situation being described.

Anderson and Ortony (1975) presented one of the earliest demonstrations of the influence of linguistic context on retrieval of object-based information. They noted that when people read sentences such as “Pianos can be pleasing to listen to,” knowledge about pianos as musical instruments was highly accessible, whereas knowledge irrelevant to the situation (e.g., that pianos are very heavy) was not. In contrast, when a sentence such as “Pianos can

be difficult to move” is read, knowledge of the physical characteristics of pianos (e.g., their weight) was highly accessible, whereas knowledge that pianos are musical instruments was backgrounded. Similar demonstrations are provided by Barsalou (1982, 1987), McCloskey and Glucksberg (1978, 1979), and Roth and Shoben (1983). Thus, it is clear that linguistic context can differentially highlight information about objects and categories.

Various explanations of the contextual nature of knowledge retrieval in these tasks have been proposed. For example, Roth and Shoben (1983) discussed accounting for context effects using semantic memory models (see, e.g., Collins & Loftus, 1975), models based on bridging inferences (see Clark, 1975), and exemplar models. However, they conclude that each account is too narrow to adequately address the complexity of the phenomena, and they suggest that an adequate proposal might draw from all three accounts. Perhaps a simpler approach to context effects is provided by embodied approaches to cognition (see, e.g., Barsalou, 1999a; Glenberg, 1997). Embodied approaches are based on two premises. First, knowledge is rooted in perceptual representations that are acquired through encounters with particular objects in particular situations. Second, the primary functions of perceptually grounded knowledge is to guide action in the world (Glenberg, 1997) and to simulate objects and situations in their absence (Barsalou, 1999a). Thus, when a sentence describes moving a piano, the simulation focuses on the actions involved in lifting and pushing the instrument but not the experience of hearing the instrument.

The claim that perception-like information is available when words are processed has been supported by a num-

This research was conducted while A.M.B. was a visiting scholar at the University of Wisconsin, Madison, supported by a Marco Polo Grant from the University of Bologna. A.M.G. received support from a University of Wisconsin Vilas Associate Award and NSF Grants BCS-0315434 and INT-0233175. We are grateful to the students who participated in the experiments and to Larry Barsalou for discussions on these topics. Thanks also to Art Graesser, Gabriel Radvansky, and two anonymous reviewers for comments on earlier drafts. Requests for reprints should be directed to A. M. Borghi, Department of Psychology, University of Bologna, Bologna, Italy, or to A. M. Glenberg, University of Wisconsin, 1202 West Johnson St., Madison, WI 53706 (e-mail: borghi@psibo.unibo.it; glenberg@wisc.edu).

ber of recent studies. For example, Pecher, Zeelenberg, and Barsalou (2003) asked people to verify that concepts had particular properties—for example, that a BLENDER is *loud*. The decision was facilitated when the previous verification trial tapped the same perceptual dimension (LEAVES—*rustling*) than when the previous trial tapped a different perceptual dimension (CRANBERRIES—*tart*). Apparently, thinking of the meaning of these concepts made perceptual information available, and, just as there is a processing cost to shifting perceptual dimensions that are being attended, there is also a cost to switching conceptual dimensions. Another finding was reported by Zwaan and Yaxley (2003), who asked participants to judge whether or not two concepts presented on successive lines in the center of a computer monitor (e.g., HEAD and FOOT) were related. Responding was faster when the words were spatially iconic (e.g., HEAD above FOOT) with their locations on the object to which they were related. Again, it appears that thinking of the meaning makes perceptual information available.

Glenberg and Kaschak (2002) present evidence that motoric information is available during language comprehension. They asked participants to verify that sentences such as “Courtney handed you the pizza” were sensible in contrast to sentences such as “Courtney handed the pizza you.” Although irrelevant to the judgment, half of the sensible sentences described actions toward the participant and half described actions away from the participant (e.g., “You handed Courtney the pizza”). In one condition, the participants responded “sensible” by moving to a response key requiring an action away from the body, and in another condition they responded “sensible” by moving to a response key requiring action toward the body. The surprising finding was that literal direction of responding interacted with the implied direction of movement in the sentence. That is, simply considering the meaning of the sentence seemed to bring to mind spatial or functional information that affects the motor system (or vice versa).

The present work is designed to extend the perceptual simulation account of the flexibility of knowledge retrieval during language processing. On the basis of the claim that the function of the perceptual representations accessed during language processing is to simulate real action in the world, the following prediction can be made: Retrieval of information about object parts depends on their spatial organization and on the functional perspective from which the information is accessed. This prediction is tested in three experiments. Experiments 1 and 2 together demonstrate that fine-grained spatial information is made available when words are used to tap conceptual information (see Murphy, 2002, for a discussion of the relationship between concepts and word meanings). Experiment 1 sets up the demonstration by illustrating an effect of functional perspective. For example, when thinking about a car, what one thinks about depends on whether one takes the perspective of driving the car or fueling the car. In Experiment 2, we examined ef-

fects of spatial organization within a perspective. Thus, from the perspective of fueling the car, information about the trunk (close to the fueling activity) is more available than information about the headlight (far from the fueling activity). From the perspective of driving the car, however, information about headlight and trunk are equally available. Experiment 3 demonstrates that this fine-grained spatial information is likely to be in the service of action. That is, thinking about the wheels of a car prepares us to act in a downward direction, whereas thinking about the roof of a car prepares us to act in an upward direction.

EXPERIMENT 1

On each trial, participants read a sentence describing an object or a location from an inside (e.g., “You are eating in a restaurant”), an outside (e.g., “You are waiting outside a restaurant”), or a mixed (e.g., “You are walking toward and entering a restaurant”) perspective. This was followed by a probe word that named either a part of the object or an associate that was not a part of the object. The participants’ task was to verify if the probe named a part of the object. There were two types of parts: those typically found inside (e.g., table) and those typically found outside (e.g., sign) the object. If the linguistically described perspective affects availability of conceptual information, then there should be a perspective (inside vs. outside) \times part location (inside vs. outside) interaction. That is, from the inside perspective it should be easier to verify an inside part, but from the outside perspective it should be easier to verify an outside part.

Method

Participants. Nineteen students at the University of Wisconsin at Madison volunteered and were either paid or given credit for research participation in their introduction to psychology classes. All the participants were native English speakers and right-handed.

Materials. Ten sets of stimuli were constructed around 10 concepts (airplane, camper, car, castle, church, elementary school, factory, house, restaurant, and train). Each set included one sentence evoking an inside (e.g., “You are driving a car”), one sentence evoking an outside (e.g., “You are washing a car”), and one sentence evoking a mixed (e.g., “You are walking toward and entering a car”) perspective, four inside probes that named parts found inside the object (e.g., steering wheel, horn, fuel gauge, gas pedal), four outside probes that named parts found outside the object (e.g., trunk, tires, antenna, door handle), and four words associated with the concept (e.g., garage, road, taxi, street) but that are not parts of the object. The latter four probes should elicit the answer “no” on the part verification test. We chose associated but nonpart probes so that the participants could not respond on the basis of a quick associative check. In addition, there were six yes/no questions for each concept. For each of the three perspectives, there was one question to which the correct answer was “no” and one for which the correct answer was “yes.” Sample questions for the car concept are “Can you see outside?” and “Can you touch the headlights?” (questions requiring positive and negative answers, respectively, for the inside perspective); “Is the car in front of you?” and “Is the car behind you?” (questions requiring positive and negative answers, respectively, for the outside perspective); and “Is the car close to you?” and “Is the car far away from you?” (questions requiring positive and negative answers, respectively, for the mixed perspective).

These questions were used to ensure that the participants paid attention to the perspective sentences and did not focus solely on the concept name. Over the course of the experiment, each concept name appeared in 42 trials. Fourteen trials began with the inside perspective sentence. Twelve of these were followed once each by the four inside probes, the four outside probes, and the four associated probes. On the other 2 trials, the perspective sentence was followed by a yes/no question. Similarly, there were 14 trials that began with the outside perspective sentence and 14 trials that began with the mixed perspective sentence. These 42 trials were equally and randomly (for each participant) divided among six blocks. Given that there were 10 concepts in total, each block consisted of 70 trials and the whole experiment consisted of 420 trials. The materials for the three experiments can be found on the Web site <http://gral.ip.rm.cnr.it/borghini/BGK04materials.htm>.

Procedure. After signing consent forms, the participants read instructions from the computer screen. The instructions indicated that a sentence would be presented on the screen until the participant depressed the space bar on the computer keyboard. Next, the word PART or QUESTION would appear for 500 msec and would be followed by a part or a question, respectively. The participants were instructed to respond “yes” or “no” as quickly as possible by using specially labeled keys on the keyboard. The computer recorded the response and the reaction time (RT). Following a series of practice trials, the 420 trials were presented. The experiment lasted 30 to 40 min.

Results

Analyses on the RTs were performed only for trials in which a probe required a “yes” response. RTs exceeding 4,000 msec (0.38% of the data) were eliminated as outliers. All analyses were conducted using a Type I error rate of .05.

Analyses of the errors revealed no evidence of a speed–accuracy tradeoff. Consequently, we focus on the RT data presented in Figure 1. For the participants analyses (indicated by F_1), condition means were obtained by averaging across concepts, and for the materials analyses (indicated by F_2) they were obtained by averaging across participants. The expected interaction between perspective location and part location was significant [$F_1(2,36) = 3.72$, $MS_e = 5,393$; $F_2(2,18) = 4.17$, $MS_e = 3,275$].

Considering just the outside perspective, the RT difference between the inside and outside parts was not significant [$F_1(1,18) = 2.45$, $MS_e = 5,483$, $p = .13$; $F_2(1,9) = 1.69$, $MS_e = 7,136$, $p = .23$]. Considering only the inside perspective, the RT difference between inside and outside parts was significant [$F_1(1,18) = 7.05$, $MS_e = 3,879$; $F_2(1,9) = 4.9$, $MS_e = 2,916$, $p = .054$].

Discussion

Much in accordance with Anderson and Ortony (1975), Roth and Shoben (1983), and others, our results show an effect of perspective on the availability of conceptual knowledge. These results extend the phenomenon in two ways. First, the perspective effect concerns object parts rather than other attributes (e.g., that pianos are heavy or sound good). Second, the perspective effect seems to involve a representation of spatial or functional location of parts. For example, the steering wheel of a car is conceptualized as close when one is in the car and as farther away when one is outside the car. Importantly, this experiment sets the stage for demonstrating an exquisite type of flexibility—namely, that perspective controls not just access to parts but information about the spatial/functional relations among parts.

It should be noted that the comprehension questions (in addition to the stimulus sentence descriptions) may have encouraged the participants to monitor perspective. We have reasons to believe, however, that these effects would occur even in the absence of an orienting task that emphasized the construction of a perspective. Recent evidence suggests that perspective is automatically produced as part of the mental simulation (for a thorough review on perspective, see MacWhinney, in press, and Zwaan & Madden, in press).

EXPERIMENT 2

Consider the parts of a car from the perspective of filling the gas tank: The trunk of the car will be closer than

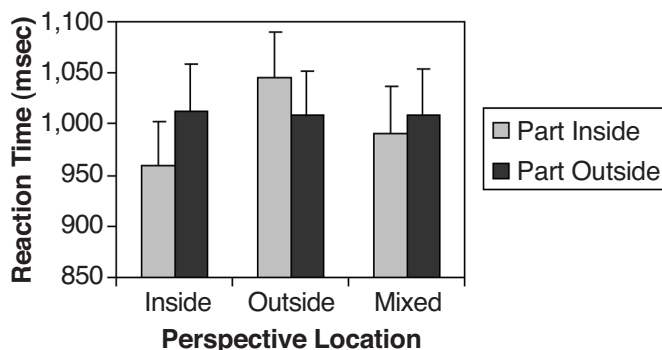


Figure 1. Data from Experiment 1: mean reaction time to determine that a probe word named a part of an object as a function of perspective on the object and location of the part. The error bars represent 1 standard error.

the headlight. Thus, if fine-grained spatial information about a concept is made differentially available by a linguistically suggested perspective, from the perspective of fueling a car it should be easier to verify that a trunk is a part of that car than that a headlight is a part of the car.

Consider further the possibility that functional perspective controls the relevance (and thus the accessibility) of information about parts. For example, from the outside perspective of fueling a car, the distance of inside parts such as the steering wheel and the back seat are irrelevant because interaction is easy with neither of them. Similarly, from the inside perspective of driving a car, the distance between outside parts such as the trunk and the headlight is irrelevant. Consequently, in verifying that an object has particular parts, there may be a three-factor interaction among perspective (inside vs. outside), part location (inside vs. outside), and part distance (near vs. far). That is, part distance should affect part verification only for those parts whose locations are functionally relevant from a particular perspective.

Method

Participants. Seventeen students at the University of Wisconsin at Madison were paid for their participation. All the participants were native English speakers and right-handed.

Materials. The materials were based on the 10 concepts used in Experiment 1. Each set of concepts contained an inside perspective sentence and an outside perspective sentence, but we did not use

the mixed perspective sentences. There were four inside parts and four outside parts, but these were subdivided into two sets of two: near parts and far parts. *Near* and *far* were defined in terms of likelihood of interaction given the perspective sentences. As in Experiment 1, there were four associated nonparts and six yes/no questions. Other than the facts that there were no mixed perspective sentences and that the parts were classified as near and far, the structure of the experiment was identical to that of Experiment 1.

Results

The most important data are presented in Figure 2. Once again, there was no evidence of a speed–accuracy tradeoff, although 2 participants were eliminated for making more than 20% errors. RT analyses were performed only on correct trials. Also, RTs longer than 4,000 msec, corresponding to 0.46% of the correct data, were eliminated. For the participants analysis, condition means were obtained by averaging across concepts, and for the materials analysis they were obtained by averaging across participants. There was a significant three-factor interaction involving perspective (inside vs. outside), part location (inside vs. outside), and distance (near vs. far) [$F_1(1,14) = 6.35$, $MS_e = 5,039$; $F_2(1,9) = 5.13$, $MS_e = 4,251$]. Thus, for the inside perspective, the participants responded to inside near parts more quickly than to inside far parts, but distance had little effect for outside parts. In contrast, for the outside perspective, the participants responded to outside near parts more quickly

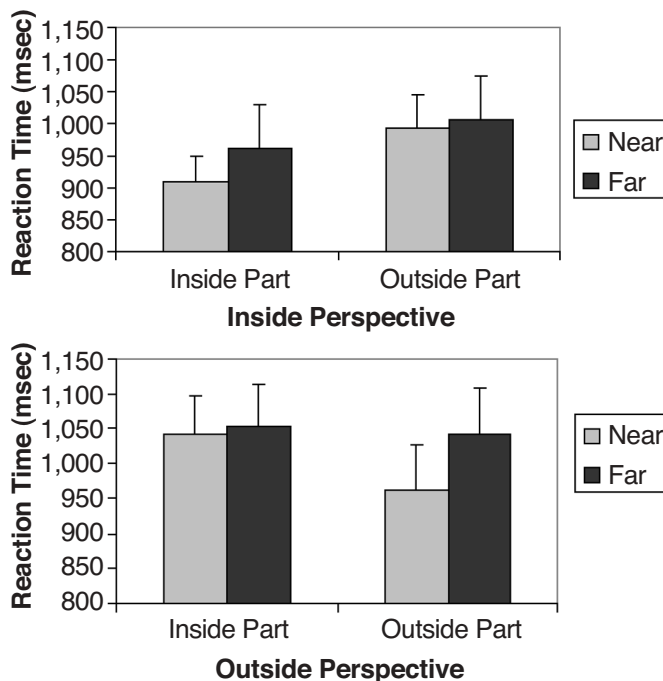


Figure 2. Data from Experiment 2: mean reaction time to determine that a probe word named a part of an object as a function of perspective on the object (top panel, inside perspective; bottom panel, outside perspective), location of the part, and distance. The error bars represent 1 standard error.

than to outside far parts, but distance had little effect for inside parts.

Several other effects were also significant. There was a main effect of perspective because responding following the inside perspective sentence ($M = 968$ msec) was faster than responding following the outside perspective sentence ($M = 1,024$ msec) [$F_1(1,14) = 24.72$, $MS_e = 3,876$; $F_2(1,9) = 4.02$, $MS_e = 13,684$, $p = .08$]. As was the case in Experiment 1, the two-way interaction between perspective and part location was significant [$F_1(1,14) = 7.21$, $MS_e = 11,730$; $F_2(1,9) = 19.48$, $MS_e = 2,956$].

Discussion

The results replicate and extend a major finding from Experiment 1. The replication consists in the fact that there is an interaction between perspective (inside vs. outside) and part location (inside vs. outside) such that the participants are faster to verify inside parts following the inside perspective sentence and faster to verify outside parts following the outside perspective sentence. Apparently, accessibility of conceptual information is affected by linguistically conveyed perspective. The extension from Experiment 1 lies in the fact that even within a perspective there is a type of distance effect. That is, within a perspective parts that are closer (or more likely to be functionally relevant) are verified more quickly than are parts that are likely to be spatially more distant (or less likely to be functionally relevant).

An alternative explanation of the distance effect is that it reflects associative relations based simply on contiguity of experience rather than spatial/functional relations. For example, the near parts (e.g., tire and trunk) may be more highly associated with the corresponding activity (e.g., fueling a car) than the far parts (e.g., headlights and antenna) because of experience reading sentences such as "While fueling the car, she noticed that the tire was partially deflated." It is difficult to test this alternative by collecting association norms. For example, if one were to ask participants to free-associate to a verb such as *fueling*, they may well instantiate a car-fueling scenario and use the spatial or functional relations in the scenario to generate responses such as "tire." Thus, although the data might be labeled in terms of an "association," they would in fact represent spatial or functional relations. Instead, we reasoned that if the distance effect reflected some sort of associative effect, then the effect should correlate with an objective measure of co-occurrence in related contexts. That is, do the concepts *fueling* and *tire* occur in similar contexts more frequently than do *fueling* and *antenna*? To obtain an objective measure of this sort of co-occurrence, we used the latent semantic analysis (LSA) program. LSA (Landauer & Dumais, 1997) uses a high-dimensional space to measure the association between words found in texts. The space is formed by analyzing the occurrences of words in several thousand texts. Each text corresponds to a column of a matrix, and each word found in the texts corresponds to a row. The entries into

the matrix are the number of times a particular word occurs in a particular text. After some preprocessing, the dimensionality of the matrix is reduced using a mathematical algorithm similar to those used in principal components analysis, and the end result is often a matrix with 300–400 dimensions. Each word is given a value on each dimension, resulting in a vector corresponding to each word. Contextual similarity between words is measured by the cosine of the angle between two word vectors. The similarity between a sentence and a word is calculated by adding the vectors of the words in the sentence and computing the cosine between that vector and the vector for a word. Thus, these LSA vectors provide an index of a type of co-occurrence of words in similar texts, and this index has been demonstrated to be significantly correlated with psychological effects depending on associations such as semantic priming.

One trick to getting usable data from LSA is to choose the right set of texts, or *semantic space*. For example, in judging the similarity of texts written by novices to texts written by experts, the cosines should be computed using a matrix formed by sampling texts from the domain of expertise. In that way, technical uses of words and technical relations are likely to be well represented. The concepts and perspectives used in Experiment 2 refer to quite ordinary and everyday interactions such as fueling a car, eating in a restaurant, and entering a school building. Consequently, we used the matrix formed from texts characterized as "General reading up to first year of college" (see <http://lsa.colorado.edu/>). We computed the cosines between each perspective sentence and the corresponding near parts and far parts. We then used a paired t test to determine if the differences between the cosines were significant. In fact, none were. For the inside perspective sentences and the inside parts, the mean cosines for the near and far parts were .21 and .16, respectively [$t(18) = 1.37$; there were 18 degrees of freedom, because one cosine could not be determined because the corresponding word did not occur in the corpus]. The corresponding data from the inside perspective sentences and outside parts were .16 and .15, respectively [$t(19) = .17$]; those from the outside perspective sentences and inside parts were .13 and .15, respectively [$t(18) = 1.08$]; and those from the outside perspective sentences and outside parts were .09 and .13 [$t(19) = -.86$]. Hence, there is little reason to suspect that the distance effects reflect differential association based on frequency of occurrence in similar contexts.

These effects of spatial distance are reminiscent of Kosslyn, Ball, and Reiser's (1978) work on imagery. The basic idea proposed by Kosslyn et al. was that when attention is focused on one region of an image and then is shifted to another region, the shift involves a scan across the space represented in the image. Furthermore, if scanning takes place at a constant rate, shifts between regions that are further apart should take longer, and that is exactly what they found. In some ways, our results are

comparable; that is not surprising given that both research projects were directed at uncovering principles of conceptual representation. There is, however, an important difference between the research projects. Kosslyn et al. were explicitly investigating imagery and instructed their participants to use imaginal processes. Our participants were simply told to verify whether or not objects had particular parts. Thus, as was noted by Barsalou, Solomon, and Wu (1999), the processes underlying verbal access to conceptual information often mirror the processes used in imagery tasks, as if the conceptual information had a perceptual base.

A similar idea is proposed in work on deixis by Duchan, Bruder, and Hewitt (1995). By their account, readers of a new text construct a deictic center (DC)—that is, an organized representational system that combines information from the text with the reader's world knowledge. The DC allows readers to gather and maintain information about who the participants are and where and when the events take place in the narration. In particular, the authors claim that the protagonist of the narrative is followed by a mental camera, which serves as a mind's eye. The camera has a focus and an angular perspective, is fuzzy around the fringes of the focus, and moves across the terrain. Information is most activated in the focus, followed by the fuzzy fringes, followed by a distant locus (because it takes time to move the camera). The idea of the mind's eye readily explains the distance effects we found in accessing parts and spatial relations.

Several questions remain unanswered. Among them is whether the results reflect differential access to knowledge because of spatial relations (e.g., while fueling a car the headlight is at a greater distance than the trunk from the location where the activity takes place) or functional possibilities (e.g., while fueling a car it is more difficult to interact with the headlights than with the trunk). Another question is whether the effects arise because of the explicit actions mentioned in the setting sentences. That is, is it necessary for the participant to consider an activity such as driving or fueling for the distance effects to arise, or is accessibility of object parts affected by spatial and/or functional distances without an explicit activity context? A final question is in regard to the nature of the retrieved information: Should the information be characterized as symbolic, spatial, or based on perception/action codes as suggested by Glenberg and Kaschak (2002) and by Hommel, Müseler, Aschersleben, and Prinz (2001)? Experiment 3 was designed to begin to answer these questions.

EXPERIMENT 3

Tucker and Ellis (1998) provided some evidence that action features are part of conceptual representations. In one experiment, they presented participants with photographs of common objects that have parts conventionally used for interaction, such as a cup with a handle. The

task was to judge if the object was upright or upside down by making a buttonpress with the left or right hand. Tucker and Ellis (1998) observed an interaction such that correct responding was facilitated when the hand assigned to the correct response was on the same side as the part used for interacting with the object. For example, suppose that a picture of a right-side-up cup was presented and the right hand was assigned to the "yes" (upright) response. In this case, responding "yes" was faster if the handle of the cup was depicted on the right (same hand as the correct response) than if it was depicted on the left. These results might be explained by the automatic elicitation of affordances. An *affordance* (Gibson, 1979) of an object is defined by how it supports interactions of particular types, such as picking up a cup. Whether or not an object affords a particular type of interaction is determined jointly by the physics of objects (e.g., the size, shape, and rigidity of the handle) and the biology of an animal (e.g., the size, strength, and agility of fingers). Thus, for example, a cup handle affords hefting by a human, but not by a dog. In the Tucker and Ellis (1998) situation, the picture of the cup may induce the perceptual system to derive affordances that facilitate interaction with the object and, in this case, a compatible response.

Ellis and Tucker (2000) and Tucker and Ellis (2001) further demonstrated the effects of affordances by presenting participants with real objects that were categorized as natural or artifact. In one set of experiments, the category response was indicated by the direction of wrist rotation (e.g., a clockwise rotation corresponds to the classification as artificial). Tucker and Ellis (2001) found a compatibility effect such that when the correct response was indicated by the same sort of wrist motion as that needed to grab the object, responding was quicker than when the two motions were mismatched. For example, consider responding to a plastic bottle or to a toothbrush when a clockwise wrist rotation indicates "artifact." Grasping a bottle usually requires a clockwise wrist rotation, whereas grasping a toothbrush (lying on a counter) usually requires a counterclockwise rotation. In this case, responding "artifact" for the bottle was faster when "artifact" corresponded to a clockwise rotation, whereas responding "artifact" for the toothbrush was faster when "artifact" corresponded to a counterclockwise rotation. Thus, with actual objects, access to conceptual information (i.e., natural or artifact) was affected by functional relations. Previous evidence showing a linkage between language (in contrast to pictures and objects) and action has been provided by Glenberg and Kaschak (2002), as was described in the introduction.

In Experiment 3, we investigated whether conceptual knowledge tapped by linguistic stimuli exhibits sensitivity to spatial relations and potential action even when the language does not mention or imply action and does not provide a functional perspective. Thus, Experiment 3 is similar to the experiment of Tucker and Ellis (1998) in that there is no explicit requirement or suggestion to con-

sider action, but it differs from Tucker and Ellis (1998) in using linguistic stimuli to tap conceptual knowledge instead of pictures or objects. Experiment 3 is similar to the experiment of Glenberg and Kaschak (2002) in the use of linguistic stimuli (and a similar response apparatus), but it differs from that of Glenberg and Kaschak in that there are no explicit action contexts.

In Experiment 3, the participants used a button box oriented vertically so that three buttons were arranged from top to bottom. The button box was mounted between the participant and the computer screen. A perspective sentence appeared on the computer screen. The sentence described an object in a particular orientation but did not describe any action (e.g., "There is a car in front of you"). In the movement condition, after reading a perspective sentence the participant used his or her right hand to push and hold the central button to see a part name. The part name was visible only while the central button was being depressed. The participant's task was to verify that the object named in the perspective sentence included the part. In the yes-is-up condition, the participants moved the right hand to an upper button to indicate "yes" and to a lower button to indicate "no." In the yes-is-down condition, the response assignment was reversed. Orthogonally, half of the parts (the upper parts) were near the top of the object (e.g., roof), and half (the lower parts) were near the bottom (e.g., wheel).

Suppose that contacting conceptual information from a linguistic cue helps access spatial information (such as the location of the object's parts) and action information (such as how to interact with the object and its parts), just as Tucker and Ellis (1998) demonstrated for picture cues. For example, simply noting that a car has a roof prepares one to interact with that part, by moving up. This preparation should facilitate a "yes" response in the yes-is-up condition and interfere with a "yes" response in the yes-is-down condition. Likewise, noting that a car has wheels prepares one to interact in a downward manner by facilitating a "yes" response in the yes-is-down condition and interfering with a "yes" response in the yes-is-up condition. This interaction between actual response direction and conceptual knowledge is similar to what Glenberg and Kaschak (2002) referred to as an *action-sentence compatibility effect*, or ACE.

In the no-movement condition, the participant kept the index finger of his or her right hand on the Yes button and the index finger of the left hand on the No button. Both the perspective sentence and the probe were presented without the participant's intervention. Thus, responding did not require large arm movements. If the linguistic cue automatically accesses spatial information, then simply responding "yes" by pressing the yes-is-up button should be fast for upper parts and slow for lower parts.

Thus, the design of Experiment 3 allows us to answer several questions. First, if spatial and/or action information is accessed through a linguistic cue even without explicit mention of action, then we should observe an ACE

effect. Second, we can determine if the information retrieved is solely spatial or includes an action (and thus a functional) component. Suppose that ACE effect is observed in the no-movement condition. That would imply that an incongruity between spatial information (e.g., that wheels are low on a car) and response location (e.g., yes is up) is sufficient to produce an ACE effect. In contrast, suppose that ACE effect is observed in the movement condition but not in the no-movement condition. This finding would imply that an incongruity between spatial information and response location is not sufficient to produce an ACE effect. The effect instead requires real action (as in the movement condition), implying that the incongruity is between real action and action information (e.g., affordances) retrieved from the concept, not just spatial information.

Method

Participants. Twenty-one students at the University of Wisconsin at Madison were paid to participate. The first 9 students were assigned to the movement condition, and the final 12 were assigned to the no-movement condition. Within each movement condition, approximately half of the participants were assigned to the yes-is-up condition for the first half of the experiment and to the yes-is-down condition for the second half. This assignment was reversed for the remaining participants.

Materials. The materials consisted of 42 concept sets. Each set consisted of a sentence such as "There's a doll standing on the table in front of you." Because it is not clear whether some objects have canonical orientations (Palmer, Rosch, & Chase, 1981), the sentences provided an orientation but did not suggest any action on the part of the participant. Each set also contained four probe words. Two of the probe words were parts of the object named by the subject noun, one of which was an upper part (e.g., hair) and one of which was a lower part (e.g., ankle). The other two probe words were associated words that were not parts of the object named by the subject noun (e.g., *kindergarten* and *baby*). During the experiment, each sentence was presented four times and was followed by a different probe word each time. As in Keil (1989), 14 of the subject noun concepts were small artifacts (e.g., doll, candle, piano), 14 were living things (e.g., dog, tree, cornstalk), and 14 were large objects and locations (e.g., mountain, skyscraper, forest).

Procedure. After signing consent forms, the participants read instructions from the computer screen. For the first 9 participants (who were in the movement condition), the instructions explained the part verification task and the yes-is-up or the yes-is-down condition. On each trial, the perspective sentence was presented for a duration dependent on sentence length—that is, the total duration was 1,000 msec plus 20 msec for each character in the sentence. Following presentation of the sentence, the computer waited for the participant to press the central button. While that button was depressed, the probe word was presented. The participants had been instructed to move to the Yes or No button as quickly as possible after deciding if the probe word named a part of the object. There were 12 practice trials and 84 critical trials. The critical trials consisted of a random ordering of 2 trials presenting each of the 42 concepts, of which 1 trial was a yes trial (randomly assigned either to an upper part or to a lower part) and 1 trial was a no trial. Following these trials, the participants were further instructed about switching the labels on the response keys so that those who began in the yes-is-up condition moved to the yes-is-down condition and vice versa. The change in response location was followed by 12 additional practice trials and the remaining 84 critical trials.

The no-movement condition differed from the movement condition in two respects. First, the participants kept their hands on the Yes (right hand) and No (left hand) buttons. Second, after a delay of 250 msec following presentation of the perspective, the probe word was presented automatically and stayed on the computer screen until the participant made a response.

Results

The data from 3 of the first 9 participants in the no-movement condition were discarded because the participants committed more than 20 errors. These 3 participants were replaced so that there was a total of 9 participants contributing data to each of the movement and no-movement conditions. An analysis of the errors revealed no evidence of a speed–accuracy tradeoff. For the RTs, only correct trials were analyzed. In addition, we eliminated RTs exceeding 4,000 msec (1.98% of the correct trials in the movement condition and 1.28% of the correct trials in the no-movement condition). Again, both the analyses on participants and those on materials were performed. The data of main interest are presented in Figure 3.

The most important finding was a significant three-factor interaction involving movement condition, part location, and response location [$F_1(1,16) = 12.11$, $MS_e = 32,510$; $F_2(1,39) = 3.69$, $MS_e = 110,538$, $p = .06$].

Consider first the movement condition. For the upper parts, responding was 124 msec faster for yes-is-up than for yes-is-down. In contrast, for the lower parts, responding was 121 msec slower for yes-is-up than for yes-is-down, resulting in a 245-msec interaction effect, which

was significant [$F_1(1,8) = 34.76$, $MS_e = 11,647$]. This interaction was not significant in the materials analysis [$F_2(1,41) = 2.09$, $MS_e = 115,299$, $p = .15$]. Examination of the stimuli revealed that one particularly uninformative stimulus item (i.e., “You are in an office”) might have adversely affected the materials analysis. When that item was removed, the effect became significant [$F_2(1,40) = 4.65$, $MS_e = 90,284$, $p = .037$]. Thus, there was an ACE effect for the movement condition.

In the no-movement condition, for the upper parts responding was 72 msec slower for yes-is-up than for yes-is-down. For the lower parts, responding was 25 msec faster for yes-is-up than for yes-is-down. This 97-msec effect is in the direction opposite that of the movement condition and is not significant [$F_1(1,8) = .60$, $MS_e = 53,909$; $F_2(1,41) = 1.66$, $MS_e = 93,861$, $p = .20$]. This pattern, which indicates an ACE effect for the movement condition and its absence for the no-movement condition, implicates action information, such as affordances, as opposed to spatial knowledge alone, as underlying the effects.

We also found a main effect of movement [$F_1(1,16) = 7.96$, $MS_e = 1,223,928$; $F_2(1,39) = 418.63$, $MS_e = 67,989.88$] due to the fact that RTs in the movement condition were longer than those in the no-movement condition. Also, the main effect of object kind was significant [$F_1(4,32) = 9.32$, $MS_e = 54,439$; $F_2(2,39) = 6.17$, $MS_e = 149,125$] due to the fact that natural kind concepts ($M = 1,159$) were processed faster than both artifact concepts ($M = 1,283$) and location concepts ($M = 1,319$). The move-

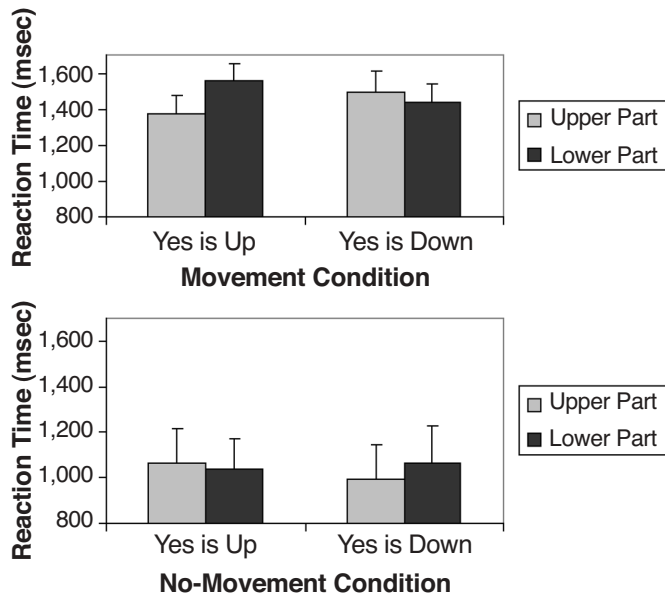


Figure 3. Data from Experiment 3: mean reaction time to determine that a probe word named a part of an object as a function of location of the object, location of the response, and whether or not a whole arm movement was required to execute the response (top panel, movement condition; bottom panel, no-movement condition). The error bars represent 1 standard error.

ment condition interacted with object kind [$F_1(2,32) = 5.06$, $MS_e = 54,439$; $F_2(1,39) = 3.40$, $MS_e = 67,990$]. This interaction resulted from the fact that the locations were slowest in the movement condition, whereas the artifacts were slowest in the no-movement condition.

Discussion

The data from Experiment 3 reinforce the conclusions drawn from Experiments 1 and 2: Conceptual information accessed from linguistic cues provides some sort of spatial information—in this case, in the up/down dimension. The experiment also provides two new results. First, the information is accessed even when the linguistic cues do not describe any action. Second, it appears that action information, or affordances, are accessed in addition to or in lieu of purely spatial information. That is, the interference between literal response location and implied part location (i.e., the ACE effect) was evident only when the participants were required to move to the response locations to make the response; there was no ACE effect when the participants simply responded in a location spatially incongruent with the retrieved information.

The results of this experiment provide an extension of Tucker and Ellis's (1998, 2001; Ellis & Tucker, 2000) demonstration that affordances are detected in photographs of objects. The ACE effect that we observed here suggests that affordances can also be derived from the information retrieved during language processing. This finding is in keeping with data showing the role of affordances in other kinds of language comprehension tasks (see, e.g., Kako & Trueswell, 2000; Kaschak & Glenberg, 2000) and lends further support to the notion that language processing (and linguistic meaning) are grounded in perception and action.

One objection to experiments that demonstrate the retrieval of perceptual or motoric information during language processing is that the demands of the task (e.g., making a movement for the task response) lead the participant to access perceptual information when he or she ordinarily would not. The results of Experiment 3 help to rule out this objection. The fact that the wheels of a car are generally located below the roof is in no logical way relevant to the task of judging whether or not wheels are a part of a car. Nonetheless, as Zwaan and Yaxley (2003) demonstrated, participants take the location of the part into account when making the judgment and responding. The conclusion is that spatial (Zwaan & Yaxley, 2003) as well as functional/action (Experiment 3) information is automatically accessed when conceptual information is used.

GENERAL DISCUSSION

Experiments 1–3 replicate the well-known finding that knowledge retrieval during language processing is context dependent (see, e.g., Anderson & Ortony, 1975) and extend this finding in several ways. First, the contextual retrieval was demonstrated for object parts rather

than for object attributes. Second, perspective effects were demonstrated even when linguistic cues do not explicitly mention action. Third, the perspective-based information seems to be centered around functional (i.e., action-based) relations such as affordances, rather than around purely spatial relations. Taken together, these findings provide additional evidence for the claim that cognition is grounded in perception and action (cf. Barsalou, 1999a; Glenberg, 1997).

Consider how the embodied approach to conceptualization addresses the results of the data presented here. The linguistic description, such as “You are fueling a car,” guides a particular simulation. During this simulation, some components of the perceptual symbol (cf. Barsalou, 1999a) of the car will be relevant to the simulation and, hence, easily accessible, and other parts will be relatively inaccessible. Which components fall into which category will depend on the individual's experiences from which the perceptual symbol has been abstracted. Thus, for most of us, “license plate” would not be particularly relevant to fueling, but for those who drive cars of a certain vintage in which the license plate is attached to the fuel tank door, “license plate” would be verified quickly. Similarly, if during fueling one spends a lot of time making faces at one's young offspring in the back seat, then “back seat” would be verified quickly. The general point is that perspective effects will depend on both individual experiences and the particular simulators (cf. Barsalou, 1999a) that are applied to the experiences.

Barsalou (2002) and Glenberg (1997) have argued that the purpose of conceptualization is to guide action. The connection between conceptualization and action can be made explicit through the concept of affordances. Recall that an affordance is support for an interaction between an animal with a particular type of body and a particular object. Thus, a small toy car affords driving by a pet hamster but not by an adult human, whereas a real car (without special apparatus) affords driving by an adult human but not by a child or a pet hamster. Glenberg and colleagues (e.g., Glenberg & Robertson, 2000; Kaschak & Glenberg, 2000) have claimed that affordances can be derived from perceptual symbols. In fact, it is claimed that these affordances are what ground the meaning of the symbols and, hence, the meaning of language (Glenberg & Kaschak, 2002).

In the discussion of Experiment 2, we noted the similarity between our results and those of Kosslyn et al. (1978). Kosslyn et al. found distance effects in an imagery task similar to the distance effects noted here. An important difference between their work and ours is that we never asked our participants to engage in imagery; instead, they simply verified whether or not objects had particular parts. The similarity of the findings in the two studies supports the proposal that conceptual representation of objects includes perceptual information (see Barsalou et al.'s [1999] discussion of instructional equivalence).

Although we interpret these findings in terms of an embodied theory of cognition, other accounts of the re-

sults can be developed via theories that eschew perceptual information in the representation of knowledge. Of particular interest here is a class of theories that may be termed *abstract, arbitrary, amodal* (AAA) *symbol theories*. With this alternative, knowledge is viewed as stored in symbolic representations that are amodal (i.e., they are descriptions rather than codes of perceptual information), abstract, and arbitrary (i.e., there is no principled relationship between the symbol and the thing being represented). The most successful AAA models incorporate two levels of representation: a purely symbolic layer that allows for the performance of cognitive operations, and a procedural layer that operates on the symbols and informs the symbolic layer on how to operate in the world.

This kind of AAA theory could, in principle, account for our data in the following manner. Suppose that the representation of a car was akin to a structural description (see, e.g., Norman & Rumelhart, 1975; Palmer, 1977). In such a representation, the spatial relations between the parts are described. For example, the symbol representing the roof of a car might be connected to symbols representing the windows, and the symbol representing the wheels would be connected to the symbol representing the axles. Furthermore, some procedures might be developed for interrogating the description (e.g., by following links between the symbols) and other procedures for using the description to guide action (e.g., when examining the wheels, reach down). Thus, when asked if “trunk” is part of “car” from the perspective of fueling, the answer comes quickly because in the structural description there are fewer links between “fuel tank door” and “trunk” than between “fuel tank door” and “headlight.” Furthermore, when asked if “wheel” names a part of a car, the procedures for interacting with a wheel are activated, thereby facilitating the literal movement of responding in a downward direction, as in Experiment 3.

Although it may be possible to account for our data using an AAA approach, there are several reasons for which we believe it is not fruitful to do so (see Glenberg & Robertson, 2000, for additional arguments). First, AAA theories rely heavily on prestored knowledge to perform cognitive operations. As has been discussed at length elsewhere (see, e.g., Barsalou, 1999b; Dreyfus, 1992), this reliance on prestored knowledge leads to a problem of combinatorial explosion. To appreciate this explosion, first note that some parts of a car such as “back seat” are literally close to the “fuel cap door.” The reason that reading about fueling a car does not facilitate responding to “back seat” is that within the perspective of fueling a car (e.g., standing outside of the car), the back seat is functionally removed from the situation. Might the links in the structural description be functional rather than spatial? This move makes the problem worse: There are so many different functional stances in regard to a car (e.g., fueling, washing, driving, repairing) that it would be cumbersome (to say the least) to encode all of

the different functional links. That is, one set of functional links would be relevant for fueling, a different set would be relevant for driving, and a very different set for selling. In addition, all of these functional links would have to be associated with different procedures. That is, from the functional perspective of fueling, the node for “wheels” is linked to a procedure such as “reach down.” However, from the functional perspective of examining a car on a lift, the node for “wheels” would be associated with procedures for reaching up. The number of symbols, propositions, and procedures required to account for complex behavior is so prohibitively large (and perhaps infinite) that it makes the problem of developing an AAA theory that displays the flexibility of human cognition intractable. Indeed, many researchers in robotics and artificial intelligence have come to the same conclusion and have abandoned this approach to system development (Dreyfus, 1992; Nolfi & Floreano, 2000). It should be recognized, however, that in the field of robotics there are also symbolic procedural, or hybrid, models that are successful in adapting to a changing environment with a certain degree of flexibility (Arkin, 1998; Hertzberg, 1999; Matarić, 2001, 2002). However, the analysis of these models is well beyond the scope of this paper.

A second reason to favor the embodied approach over AAA theories comes from neuroscience. Several studies have shown that language processing involves the recruitment of perceptual and motor areas of the cortex (Gerlach, Law, & Paulson, 2002; Martin, Haxby, Lalonde, Wiggs, & Ungerleider, 1995; Martin, Wiggs, Ungerleider, & Haxby, 1996; Pulvermüller, 1999). For example, Martin, Ungerleider and Haxby (2001) reviewed experiments demonstrating that processing words that name instruments (e.g., *hammer, knife*) leads to activation of areas in motor cortex, just as processing vision words (such as the names of colors) activates areas of visual cortex. It appears that information from different perceptual modalities shapes processing in all regions of the cortex. On the other hand, neuroscience has yet to produce any evidence for something like an AAA symbol-processing system. Thus, the embodied approach is largely consistent with what is known about neuroanatomy, whereas the AAA symbol approach is not. Thus, whereas some AAA theories may be able to provide an explanation for the data observed in these and other experiments, we believe that the embodied approach provides a simpler and more psychologically (and neurologically) plausible way to explain the results.

In conclusion, we have produced the following new data about conceptual information available from linguistic stimuli: (1) Information about parts of objects is differentially accessible depending on the perspective (inside or outside) on the object; (2) within a perspective, information about parts is differentially available depending on both the spatial relations among the parts and the functional relations determined by the perspective; and (3) one component of the functional informa-

tion is an affordance—that is, how to reach for a part and interact with it. These findings follow from the view that conceptual information is used primarily to guide interactions with the world. Theories of conceptual representation constructed from AAA symbols can be modified to fit these particular data; however, for reasons of parsimony and plausibility, these theories are less promising accounts of the phenomena.

REFERENCES

- ANDERSON, R. C., & ORTONY, A. (1975). On putting apples into bottles: A problem of polysemy. *Cognitive Psychology*, *7*, 167-180.
- ARKIN, R. C. (1998). *Behavior-based robotics*. Cambridge, MA: MIT Press.
- BARSALOU, L. W. (1982). Context-independent and context-dependent information in concepts. *Memory & Cognition*, *10*, 82-93.
- BARSALOU, L. W. (1987). The instability of graded structure: Implications for the nature of concepts. In U. Neisser (Ed.), *Concepts and conceptual development: Ecological and intellectual factors in categorization* (pp. 101-140). Cambridge: Cambridge University Press.
- BARSALOU, L. W. (1999a). Perceptual symbol systems. *Behavioral & Brain Sciences*, *22*, 577-609.
- BARSALOU, L. W. (1999b). Author's response: Perceptual symbol systems. *Behavioral & Brain Sciences*, *22*, 637-660.
- BARSALOU, L. W. (2002). Being there conceptually: Simulating categories in preparation for situated action. In N. L. Stein, P. J. Bauer, & M. Rabinowitz (Eds.), *Representation, memory, and development: Essays in honor of Jean Mandler* (pp. 1-16). Mahwah, NJ: Erlbaum.
- BARSALOU, L. W., SOLOMON, K. O., & WU, L. L. (1999). Perceptual simulation in conceptual tasks. In M. K. Hiraga, C. Sinha, & S. Wilcox (Eds.), *Cultural, typological, and psychological perspectives in cognitive linguistics: The proceedings of the 4th conference of the International Cognitive Linguistics Association* (Vol. 3, pp. 209-228). Amsterdam: John Benjamins.
- CLARK, H. H. (1975). Bridging. In R. Schank & B. Nash-Webber (Eds.), *Theoretical issues in natural language processing* (pp. 188-193). Cambridge, MA: MIT Press.
- COLLINS, A. M., & LOFTUS, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, *82*, 407-428.
- DREYFUS, H. L. (1992). *What computers still can't do: A critique of artificial reason*. Cambridge, MA: MIT Press.
- DUCHAN, J. F., BRUDER, G. A., & HEWITT, L. A. (1995). *Deixis in narrative: A cognitive science perspective*. Hillsdale, NJ: Erlbaum.
- ELLIS, R., & TUCKER, M. (2000). Micro-affordance: The potentiation of components of action by seen objects. *British Journal of Psychology*, *91*, 451-471.
- GERLACH, C., LAW, I., & PAULSON, O. B. (2002). When action turns into words: Activation of motor-based knowledge during categorization of manipulable objects. *Journal of Cognitive Neuroscience*, *14*, 1230-1239.
- GIBSON, J. J. (1979). *The ecological approach to visual perception*. New York: Houghton Mifflin.
- GLENBERG, A. M. (1997). What memory is for. *Behavioral & Brain Sciences*, *20*, 1-55.
- GLENBERG, A. M., & KASCHAK, M. P. (2002). Grounding language in action. *Psychonomic Bulletin & Review*, *9*, 558-565.
- GLENBERG, A. M., & ROBERTSON, D. A. (2000). Symbol grounding and meaning: A comparison of high-dimensional and embodied theories of meaning. *Journal of Memory & Language*, *43*, 379-401.
- HERTZBERG, J. (1999). Planning. In J. Webster (Ed.), *Encyclopedia of electrical and electronics engineering* (Vol. 16, pp. 493-504). London: Wiley.
- HOMMEL, B., MÜSSELER, J., ASCHERSLEBEN, G., & PRINZ, W. (2001). The theory of event coding (TEC): A framework for perception and action planning. *Behavioral & Brain Sciences*, *24*, 849-937.
- KAKO, E., & TRUESWELL, J. C. (2000). Verb meanings, object affordances, and the incremental restriction of reference. *Proceedings of the 22nd Annual Conference of the Cognitive Science Society* (pp. 256-261). Hillsdale, NJ: Erlbaum.
- KASCHAK, M. P., & GLENBERG, A. M. (2000). Constructing meaning: The role of affordances and grammatical constructions in sentence comprehension. *Journal of Memory & Language*, *43*, 508-529.
- KEIL, F. C. (1989). *Concepts, kinds, and cognitive development*. Cambridge, MA: MIT Press.
- KOSSLYN, S. M., BALL, T. M., & REISER, B. J. (1978). Visual images preserve metric spatial information: Evidence from studies of image scanning. *Journal of Experimental Psychology: Human Perception & Performance*, *4*, 47-60.
- LANDAUER, T. K., & DUMAIS, S. T. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction and representation of knowledge. *Psychological Review*, *104*, 211-240.
- MACWHINNEY, B. (in press). The emergence of grammar from perspective. In D. Pecher & R. A. Zwaan (Eds.), *Grounding cognition: The role of perception and action in memory, language, and thinking*. Cambridge: Cambridge University Press.
- MARTIN, A., HAXBY, J. V., LALONDE, F. M., WIGGS, C. L., & UNGERLEIDER, L. G. (1995). Discrete cortical regions associated with knowledge of color and knowledge of action. *Science*, *270*, 102-105.
- MARTIN, A., UNGERLEIDER, L. G., & HAXBY, J. V. (2001). Category specificity and the brain: The sensory-motor model of semantic representations of objects. In M. S. Gazzaniga (Ed.), *The cognitive neurosciences* (2nd ed., pp. 1023-1036). Cambridge, MA: MIT Press.
- MARTIN, A., WIGGS, C. L., UNGERLEIDER, L. G., & HAXBY, G. V. (1996). Neural correlates of category specific knowledge. *Nature*, *379*, 649-652.
- MATARIĆ, M. J. (2001). Behavior-based robotics. In R. A. Wilson & F. C. Keil (Eds.), *The MIT encyclopedia of the cognitive sciences (MITECS)* (pp. 74-77). Cambridge, MA: MIT Press.
- MATARIĆ, M. J. (2002). Situated robotics. In *Encyclopedia of cognitive science* (Vol. 4, pp. 25-30). London: Nature Publisher Group, Macmillan Reference.
- MCCLOSKEY, M. E., & GLUCKSBERG, S. (1978). Natural categories: Well defined or fuzzy sets? *Memory & Cognition*, *6*, 462-472.
- MCCLOSKEY, M. E., & GLUCKSBERG, S. (1979). Decision processes in verifying category membership statements: Implications for models of semantic memory. *Cognitive Psychology*, *11*, 1-37.
- MURPHY, G. L. (2002). *The big book of concepts*. Cambridge, MA: MIT Press.
- NOLFI, S., & FLOREANO, D. (2000). *Evolutionary robotics: The biology, technology, and intelligence of self-organizing machines*. Cambridge, MA: MIT Press.
- NORMAN, D. A., & RUMELHART, D. (1975). *Exploration in cognition*. San Francisco: Freeman.
- PALMER, S. E. (1977). Hierarchical structure in perceptual representation. *Cognitive Psychology*, *9*, 441-474.
- PALMER, S. [E.], ROSCH, E., & CHASE, P. (1981). Canonical perspective and the perception of objects. In J. Long and A. Baddeley (Eds.), *Attention and performance IX* (pp. 135-151). Hillsdale, NJ: Erlbaum.
- PECHER, D., ZEELENBERG, R., & BARSALOU, L. W. (2003). Verifying properties from different modalities for concepts produces switching costs. *Psychological Science*, *14*, 119-124.
- PULVERMÜLLER, F. (1999). Words in the brain's language. *Behavioral & Brain Sciences*, *22*, 253-279.
- ROTH, E. M., & SHOEN, E. J. (1983). The effect of context on the structure of categories. *Cognitive Psychology*, *15*, 346-378.
- TUCKER, M., & ELLIS, R. (1998). On the relations between seen objects and components of potential actions. *Journal of Experimental Psychology: Human Perception & Performance*, *24*, 830-846.
- TUCKER, M., & ELLIS, R. (2001). The potentiation of grasp types during visual object categorization. *Visual Cognition*, *8*, 769-800.
- ZWAAN, R. A., & MADDEN, C. J. (in press). Embodied sentence comprehension. In D. Pecher & R. A. Zwaan (Eds.), *Grounding cognition: The role of perception and action in memory, language, and thinking*. Cambridge: Cambridge University Press.
- ZWAAN, R. A., & YAXLEY, R. H. (2003). Spatial iconicity affects semantic relatedness judgments. *Psychonomic Bulletin & Review*, *10*, 954-958.