

Embodied Verbal Semantics: Evidence from a lexical matching task

Shweta Narayan, Benjamin K. Bergen, Zachary Weinberg
University of California, Berkeley/International Computer Science Institute;
University of Hawai'i, Manoa; International Computer Science Institute

1 Introduction

Neural processes involved in language understanding are often not considered to be of direct interest to linguists. However, these processes place constraints on linguistic theories in that a linguistic theory that correctly predicts observable non-linguistic phenomena is generally preferable to a linguistically equivalent one that does not.

One standard way of thinking about linguistic semantics is that it is ultimately based in some kind of logical system (e.g. Barwise and Perry, 1983). Another is that it is embodied, that is, concretely based on people's interactions with the world (e.g. Fillmore, 1982; Lakoff, 1987; Fauconnier, 1997).

This second, embodied perspective makes a strong claim that the logical perspective does not; namely, that linguistic semantics is grounded in the perceptual and motor systems, and therefore reflects characteristics of these systems. We present some evidence that supports this prediction, and suggest a methodology for further investigating the question of the role of perception and action in meaning.

Recent studies have shown that some of the neural circuitry involved in the execution of motor actions is also activated when the same motor actions are seen, or when language about those motor actions is understood. (Gallese et al., 1996; Rizzolatti et al., 1996; Tettamanti et al., ms). This activation may be epiphenomenal - the simple and non-functional result of spreading excitation from perceptual to motor structures. Alternatively, this so-called mirror system might play a crucial role in our understanding of verbal semantics. In this case, at least some aspects of language understanding are crucially dependent on sensorimotor representation. If this is true, it raises serious problems for any theory in which semantics is not grounded in our sensorimotor systems.

Human premotor cortex is somatotopically organized, and has major subdivisions specialized for hand, foot and mouth actions. It has been shown in an imaging study (Buccino et al., 2001) that this somatotopic organization extends to mirror cells in premotor cortex. That is, the part of premotor cortex that becomes activated when people carry out foot actions is also activated when they watch foot actions being carried out. Similarly, a different part is activated when people

carry out or watch hand actions, and a third when they carry out or watch mouth actions. In addition, Pulvermüller et al (2001) found that verbs associated with different effectors were processed in different parts of motor cortex, and that their processing led to reaction times that varied by effector type.

If understanding a verb depends on activation of mirror circuitry, then in order to understand verbs whose meaning involves particular effectors (body parts), very specific circuits must be activated in the corresponding parts of premotor cortex. There should therefore be differential activation in premotor cortex based on effector type; and, crucially, this activation should be involved in understanding. If this is the case, we would expect that if we ask subjects to compare two different concepts involving the same effector, these two concepts will interfere with each other to a certain extent, since there is an overlap in the neural circuits required to understand them. If we compare two concepts that involve different effectors, on the other hand, there should be less interference, since there are two distinct parts of premotor cortex being activated.

We would expect increased interference rather than facilitation because the task involves comparison of similar but nonidentical concepts. Very similar concepts would have a large degree of overlap in the neural circuits that encode them; therefore, someone trying to compare two similar concepts would have two very similar neural patterns activated. These would more strongly inhibit each other since they are in close competition, and it would therefore be relatively difficult to decide that they were different.

In a previous experiment (Bergen et al., 2003), subjects were shown an image of a stick figure performing an action. This image was displayed for one second, and was followed by a 500ms interstimulus interval.¹ Subjects were then presented with a written verb. Their task was to judge, as quickly as possible, whether the verb matched the image.

The image/verb pairs fell into three categories. Matches, where the verb was a good descriptor for the image, made up half the stimuli, and provided the control condition. The other half of the stimuli were non-matches. There were two non-matching conditions: one where the image and verb involved different effectors, and one where they involved the same effector.

We hypothesized in that study that subjects would have relatively longer reaction times when comparing non-matching stimuli involving the same effector than when comparing non-matching stimuli involving different effectors, due to the greater interference between same effector concepts predicted by the embodied model. In both cases the reaction time would be higher than in the control condition (where we would not expect interference effects).

Consistent with this hypothesis; we found that mean reaction time was significantly higher for the same-effector condition than either of the other two ($p < .0001$). While mean reaction time was higher for the different-effector condition than the control condition, this difference was not significant (Table 1).

¹In the first 450ms of this, a visual mask covered the screen; this was to make sure there were no effects caused simply by visual imagery.

Embodied Verbal Semantics: lexical matching task

Table 1: Means Table for Reaction Time as an effect of Condition in the Image-Verb matching task

Condition	Count	Mean (msec)	Std. Dev.	Std. Error
Match	1596	740.57	257.54	6.45
Non-match same-effector	840	798.54	251.00	8.66
Non-match different-effector	870	750.93	204.74	6.94

However, there were several possible confounds that needed to be addressed. The first was that subjects might be reacting to some abstract semantic similarity of the particular verb-image pairs that had nothing to do with effectors. To test this possibility, we evaluated the similarity of each non-matching verb with its corresponding matching verb, using the semantic similarity metric of Latent Semantic Analysis (LSA, Landauer et al., 1998). LSA is a statistical method for judging the similarity of words based on the context in which they appear, and has been shown to perform similarly to humans in several ways, including synonym and multiple-choice tasks. LSA uses a continuous pairwise comparison function ranging from -1 to 1. That is, two words that appear in all the same contexts would have an LSA rating of 1, two words that appear in none of the same contexts would have an LSA rating of -1, and any overlap in contexts would result in a rating between -1 and 1. This analysis can therefore provide a numeric measure of abstract similarity, independent of effector type.

For each non-matching condition, we obtained the verb's LSA rating with respect to the corresponding matching verb. We then performed a regression analysis of this rating against the mean reaction time (RT) per trial. There was a very weak positive correlation between LSA rating and RT ($R = 0.094$), suggesting that more similar items tend to take longer to distinguish. However, this correlation was not significant, so while it might constitute a small, partial explanation for our results, it cannot be the full story.

Another potential confound was that subjects may have been responding to visual properties of the images rather than motor properties of the actions they depicted. A picture of one action might just resemble a subject's mental image of another action involving the same effector, regardless of whether the actions themselves were similar. If there was more visual similarity between actions that used the same effector, then this would yield greater visual interference (and thus the observed longer reaction times) when non-matching items used the same effector than when they used different ones.

In order to test this possibility, the following experiment was conducted. Subjects were shown a verb, followed by a near-synonym or a mismatch. Again the mismatches fell into two different categories; the actions being compared involved either the same effector, or a different one. Subjects were asked to

decide whether the two verbs had approximately the same meaning.² The absence of images in this design removes the possibility of a visual confound.

1.1 Hypotheses

Null hypothesis: The motor/premotor activation found when subjects read and understand action verbs is epiphenomenal, and not necessary for understanding. When subjects are asked to judge whether two verbs mean approximately the same thing, simultaneous activation of the same motor or premotor areas by both verbs will play no part in that judgment. Therefore reaction time will not be affected by whether mismatches involve the same or different effectors.

Alternate hypothesis: When subjects are presented with verb-verb pairs and asked to judge whether they match, the reaction time for mismatches will be significantly longer when the verbs share an effector than otherwise.

2 Methods

2.1 Subjects

Subjects who participated in this study were all adult native English speakers (people who had started speaking English at age 5 or younger). A few were volunteers, but most were undergraduate cognitive science students at UC Berkeley. These received either extra credit or payment for their participation; those receiving extra credit, were allowed to participate regardless of their English proficiency. However, we omitted the data from non-native speakers from our analysis. We similarly omitted data from subjects who had participated in our previous study.

The classes for which students could receive extra credit by participating in this experiment had not, at the time of the experiment, covered material that would help students to guess the study's purpose.

2.2 Design

2.2.1 General overview

Subjects were asked to sit at a computer and press labeled keys for yes or no, as appropriate. They were asked to sit comfortably and use different hands for each key, and to place their fingers on the keys before starting. At this stage, an experimenter was in the room with them.

They were then presented with written instructions. These started by informing subjects that this was a linguistics study, and reassuring them that it

² This follow-up was therefore as similar to the initial experiment as possible given the removal of images, but one change in particular led to a slight difference in the hypothesis. If the cause of the effect seen in the previous experiment really is an effect of motor activation being involved in verb understanding, we would still expect subjects to take longer to discount mismatches in the same-effector condition than the different-effector condition. But we would also predict that some interference would be involved in the comparison of near-synonyms. We would therefore predict relatively long reaction times for the matching condition. In effect, we would expect this task to be harder than matching pictures to verbs, because the matching condition does not involve exact matches.

Embodied Verbal Semantics: lexical matching task

was not an intelligence or personality test.

They were then given more detailed written instructions. These informed subjects that they would be presented with two verbs in succession, and they were to say whether they considered the two to mean “approximately the same thing”. This phrasing was necessary because it is not possible to find exact synonyms for common English action verbs. They were asked to respond to each pair as fast as possible, because there were a lot of pairs to get through; they were not explicitly informed that this was a reaction-time study.

Subjects were then asked three questions (whether they were left-handed, right-handed or ambidextrous, whether they were native English speakers, and whether they had previously participated in a related study). After this, they had a practice session, which consisted of eight randomly-chosen filler trials. During this, they were given (displayed) feedback. This part was also self-paced; subjects had to press a key to advance to the next trial, and were allowed to pause to ask the experimenter for clarification.

Once the practice session was complete, the screen informed subjects that it was, and reminded them to ask the experimenter any questions they might have. They were also reminded that there would be no feedback during the study, and that the program would continue immediately to the next pair of words once they pressed yes or no.

The experimenter then asked whether they had questions, and told them to come into the waiting area when they had completed the study. Then the experimenter left.

At the end of the study, the reminder to come into the waiting area was repeated on the computer screen, along with a message thanking subjects for their participation.

2.2.2 The Study

During the study itself, subjects were presented with a fixation cross in the center of the screen for 2 seconds.³ This was followed by an English action verb presented for 1 second, then a visual mask for 0.45 sec and a blank screen for 0.05 sec. Then the second verb was displayed, and stayed on the screen until the subject pressed "yes" or "no". At this point, the next trial started.

All verbs were capitalized and presented in the center of the screen. Verb pairs could fall into the following categories:

1. Test stimuli: These were the verb pairs we were interested in. They involved hand, foot or mouth verbs, and constituted 50% of stimuli.
2. Filler: The remaining 50% of the stimuli were fillers, which were put in to make the point of the experiment less obvious. For the fillers, we tried to pick verb pairs which did not involve hand, foot, or

³It was found that a shorter presentation of the fixation cross made the study very stressful to take, and might therefore have caused an elevated error rate.

mouth actions.⁴

In addition to this, the second verb in the pair could relate to the first in the following ways:

1. Matching: near-synonyms. 50% of the pairs presented were matching.
2. Non-matching: verbs with clearly different meanings.⁵ 50% of the pairs presented were non-matching.

The non-matching pairs for the test condition fell into two different categories:

- (a) Same effector: In this case, both verbs referred to verbs whose primary effector (hand, foot or mouth) was the same.
- (b) Different effector: the two verbs did not share a primary effector.

These were split evenly, so that 25% of the test pairs, or 12.5% of the overall pairs, fell into each of these categories. These percentages are shown in Table 2.

Table 2: Breakdown of Presentation Data

Condition	Matching	Non-matching same-effector	Non-matching different-effector	Non-matching total
Filler	25%	n/a	n/a	25%
Test	25%	12.5%	12.5%	25%

Examples of the different conditions are shown in Table 3. It should be noted that while some of the filler verbs referred to bodily action, the test condition verbs referred exclusively to human bodily action.

Table 3: Examples of Presentation Data

Condition	Initial verb	Match	Non-match same-effector	Non-match different-effector	Non-match
Test	SCREAM	SHRIEK	LICK	STEP	n/a
	TIE	KNOT	CLAP	RUN	n/a
	DANCE	WALTZ	LIMP	YELL	n/a
Filler	LEAK	DRIP	n/a	n/a	RUST
	GRAZE	FEED	n/a	n/a	PAUSE
	CHEAT	SWINDLE	n/a	n/a	BREAK

⁴In a few cases we failed, but the few students who ventured hypotheses after participating did not seem to notice the conditions we were looking at. We therefore have reason to believe that the fillers were effective.

⁵As determined by the pre-test to the previous experiment.

2.2.3 Design Issues

- ⟨ Choosing verb pairs: Each initial verb was presented twice to subjects; once followed by a near-synonym, once by a mismatch. Subjects saw each test verb with only one mismatch: either the same-effector case or the different-effector case. Therefore 50% of the correct answers were "yes". Mismatches were chosen out of the same pool of verbs as initial verbs.

Verb pairings were precalculated before the experiment to ensure that each subject saw half of the initial verbs with a same-effector non-matching verb and half with a different-effector non-matching verb. Each subject saw any given non-matching verb in the second position only once, as either a same-effector mismatch or a different-effector mismatch. Given this constraint, non-matching verbs were assigned randomly to initial verbs.

This divided subjects into two groups, depending on which set of initial verbs they saw with same-effector non-matching verbs. The survey software alternated between groups.

- ⟨ Choosing verbs: There are not many English bodily action verbs that are short, commonly-used, and unambiguous in a neutral context. As a result, we did not have a very long list of choices, and did not select based on verb frequency.⁶

Choosing near-synonyms was problematic because for the most part, there are no commonly-occurring synonyms for simple action verbs. Some of the synonyms given were arguably mismatches (for instance, WINK and BLINK). However, we omitted all 'wrong' answers from our analysis, thus using only those answers where subjects' judgments matched our own.

- ⟨ Presenting a visual mask: In the previous experiment, this was done to prevent iconic memory from affecting the results. This was much less important in the current study, but the visual mask was kept to maintain the parallel between the studies.
- ⟨ Running the experiment: Subjects were presented with a python script written by one of the experimenters, and running on a machine that was not running any other programs. The accuracy of the script's timing, when tested, came to within 10ms.

3 Results

⁶Our verbs were selected by subjects in the pretest to the previous experiment, in which they identified the verb they associated with the visual images. We chose the verbs most commonly produced by subjects in response to the visual stimuli. These verbs were used for this experiment also; this ensured that we kept a consistent stimulus set.

3.1 Data considered

Non-native English speaker data was removed from the data set. In addition, we discarded all data from subjects whose mean reaction time (RT) was more than three standard deviations from the global mean. We also removed all data from subjects whose percentage error rate was more than three standard deviations above the global mean. Items whose mean RT was more than three standard deviations from the global mean would have been removed, but there were no such items.

It is worth noting that there were some clear outliers among the subjects whose data were not eliminated from the study. Of the 53 subjects whose mean reaction time was not more than three standard deviations from the global mean, six account for every response time of more than 4 seconds, and of those, three account for every response time over 4.4 seconds. The extreme outlying points in Figure 1, therefore, are not indicative of subjects' performance in general.

For this study, we have not considered how subjects performed on the filler sentences, nor the performance of non-native English speakers.

3.2 Errors per Condition

In the previous study, the number of errors was extremely low (under 2%). There were far more errors in this study. Subjects got the correct answer for 99.1% of the different-effector mismatches, but only 93.5% of the same-effector mismatches, and 87.8% of the near-synonyms. Performance on the different-effector mismatches was significantly more accurate than performance on either the near-synonyms ($p=1.207e-10$) or the same-effector mismatches ($p=7.392e-07$), and performance on same-effector mismatches was significantly more accurate than performance on the near-synonym condition, though this difference was less pronounced ($p=1.042e-3$).⁷

3.3 Reaction Times

Counting only correct replies, the mean reaction time was 0.93 seconds for the different-effector mismatches, 1.03 seconds for the same-effector mismatches, and 1.07 seconds for the near-synonyms. Figure 1 shows the mean and quartile reaction times for the three conditions. The whiskers extend to cover all data points which are no more than 1.5 times the interquartile range from the box.

As Clark (1973) shows, standard Analysis of Variance (ANOVA) is not sufficient for any generalization beyond the subjects and linguistic items tested. In order for a further generalization to be drawn, it is necessary to determine that F' is above significance; however F' is difficult to calculate. Clark sets forth a simple method to find the minimum possible value of F' , called $\min F'$; if this is above significance then F' must also be significant.

⁷All data analysis was conducted using the *R* software package. Our thanks to Nathaniel Smith, who provided a great deal of help with both analysis and the software package.

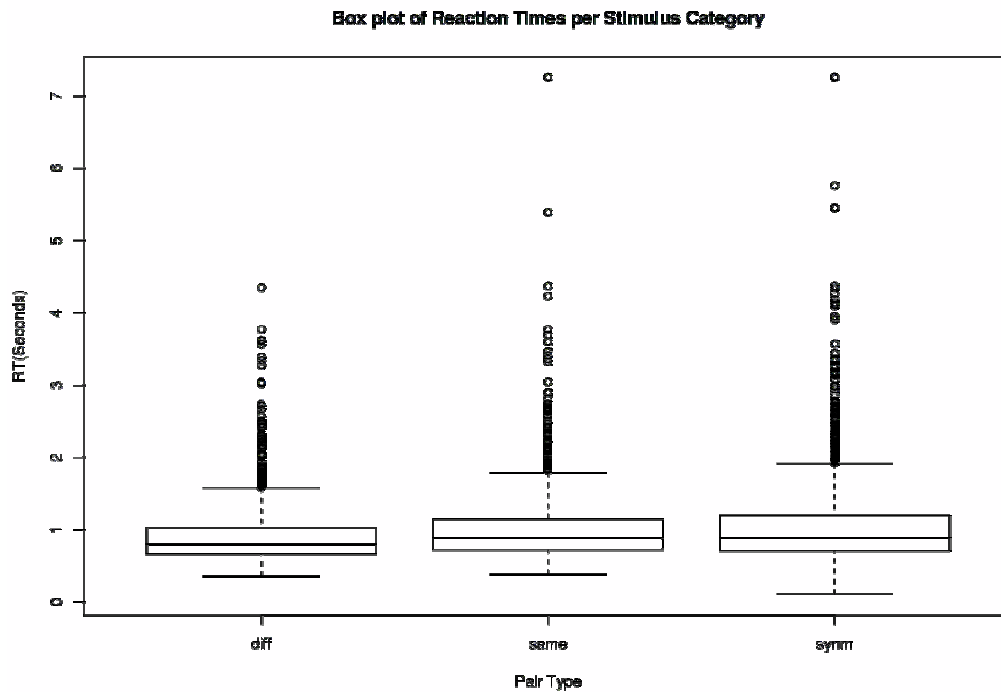
Embodied Verbal Semantics: lexical matching task

Following (Clark, 1973), we performed an ANOVA with subjects as a random factor ($F_1(1, 52) = 32.466$). We also performed an ANOVA with items as a nested random factor ($F_2(1, 82) = 12.607$). From these we determined that the difference in RT between the mismatch conditions is significant (min $F'(1, 126) = 9.0808, p = 0.0031 < 0.005$).

We therefore reject the null hypothesis.

Post-hoc tests comparing each of the non-matching conditions with the control condition give us the expected results. To compare the different-effector condition with the control, we performed an ANOVA with subjects as a random factor ($F_1(1, 52) = 25.454$). We also performed an ANOVA with items as a nested random factor ($F_2(1, 82) = 13.508$). From these we determined that the difference in RT between the control and the same-effector mismatch conditions is probably significant (min $F'(1, 126) = 9.781$). (Compare this to min $F'(1, 126) = 9.0808$ between the test conditions.)

Figure 1



To compare the same-effector condition with the control, we performed an ANOVA with subjects as a random factor ($F_1(1, 52) = 2.0015$). We also performed an ANOVA with items as a nested random factor ($F_2(1, 82) = 0.4675$). From these we determined that the difference in RT between the control and the different-effector conditions is probably not significant (min $F'(1, 129) = 0.3790$).

Since these are post-hoc tests, we cannot simply convert the min F' values to p -values; however, the values are still worth noting.

4 Discussion

Subjects took significantly longer to reject a verb pair as near-synonyms when the two verbs shared an effector than when they did not. We also know from the previous study discussed above (Bergen et al. 2003) that subjects took longer to reject a verb that did not match a visual image when the action referenced by the verb and the action depicted in the image share an effector. These results provide evidence that the meanings of motion verbs include in their representations information about bodily action, in particular, the body part that is used to perform the action.

More broadly, the effect we see in understanding studies supports the notion that linguistic semantics are *embodied*, that word meaning is closely linked to the neural circuitry used for perception and motor action. It is still possible, however, that what we have found is an effect of memory, rather than lexical semantics. That is, it is possible that verb understanding is a completely abstract process, but that holding a verb in memory long enough to compare it to another involves activation of the neural circuits discussed above. This possibility has been tested in recent work in our lab, using precisely the same methodology as the image-verb matching experiment described above, but with a reversed stimulus order, such that the written word preceded the image. The results once again demonstrated the same significant effect of shared effector on non-matching word-image pairs. The effect of memory will also be tested by a further study, in which the spoken word and image in each pair are presented simultaneously. Our hypothesis predicts that this should cause more interference if the verbs share an effector.

The work of Gallese et al. (1996) shows that in monkeys, mirror neurons can code not just for effector, but for finer distinctions, like specific handshapes. This leads to the possibility that similar actions could involve the activation of some of the same mirror neurons, in related circuits. If this is the case we might expect a fair amount of mutual inhibition between such circuits, and if so we would expect more inhibition between more similar circuits.

The work of Pulvermüller et al. (2001) and of Tettamanti et al. (ms) shows that words about motor actions cause activation of the specific mirror structures involved in performing those actions. The study presented here suggests that this activation is not epiphenomenal. When two verbs involving the same effector are presented, their neural representations are co-activated. These neural representations are similar, since they involve the same effector and therefore the same areas of premotor cortex. There is therefore a certain amount of competition between these circuits, and so it will take longer for both to become simultaneously activated, as would be necessary for the comparison required in this task. If the verbs do not share an effector, there will be less similarity between the neural representations, and they can become co-activated more

Embodied Verbal Semantics: lexical matching task

easily.⁸

MacWhinney, Glenberg, and others have argued (MacWhinney, 1999; Glenberg and Robertson, 2000; Kaschak and Glenberg, 2000), that understanding language is based on internally simulating the described scenario. For instance, when we read the sentence "Sally sauntered sedately sideways"⁹ we immediately understand the motor actions involved in sauntering; in fact, we do this so quickly that we need to reinterpret the sentence at the last word, since moving sideways would involve different motor actions from moving forward, which is our default. The results of this study are certainly consistent with these previous studies.

Regardless of the exact neural process responsible for our results, the findings reported above lend support to the hypothesis that verbal semantics are grounded in bodily experience, and our understanding of action verbs depends on our understanding of actions. This is part a growing body of work that supports the hypothesis that concepts, rather than being abstract and symbolic, are embodied in our sensorimotor experience. There is a great deal of evidence that imagined and remembered actions activate the same brain areas as performance of those actions (e.g. Wheeler et al., 2000; Nyberg et al., 2000; Nyberg et al., 2001; Porro et al., 1996; Rizzolatti et al., 1996). This study and its predecessor support the hypothesis that the same circuits are also involved in understanding language about actions. Other studies (e.g. Boroditsky et al., 2001) show that abstract domains, like time, are understood metaphorically in terms of concrete domains, like space. This suggests a way for embodied understanding of concrete concepts to map to understanding of abstract domains.

The evidence building up suggests a view of language and cognition that is ultimately grounded in our sensorimotor experience. In particular, our findings suggest that language understanding depends on the same motor circuitry involved in perceiving and producing actions, and that the organization of the sensorimotor system could be crucial to theories of language understanding.

References

- Barwise, J. and Perry, J. (1983). *Situations and Attitudes*. MIT Press, Cambridge, MA.
- Bergen, B., Narayan, S., and Feldman, J. (2003). Embodied verbal semantics: evidence from an image-verb matching task. In *Proceedings of Cognitive Science 2003*, Boston, MA.
- Boroditsky, L., Ramscar, M., and Frank, M. (2001). The roles of mind and body in abstract thought. In *Proceedings of the 23rd Annual Meeting of the*

⁸ In the case of near-synonyms, the representations share a great deal of similarity, but since they are only near-synonyms, they are not activating the same representation. Therefore there will still be mutual inhibition of circuits, and so it is not surprising that the near-synonym condition was one that showed the longest reaction times.

⁹Thanks to Alexander Lothar D'Archangel for this example.

Shweta Narayan, Benjamin K. Bergen, and Zachary Weinberg

- Cognitive Science Society*, Edinburgh, Scotland.
- Buccino, G., Binkofski, F., Fink, G. R., Fadiga, L., Fogassi, L., Gallese, V., Seitz, R. J., Zilles, K., Rizzolatti, G., and Freund, H. J. (2001). Action observation activates premotor and parietal areas in a somatotopic manner: an fMRI study. *European Journal of Neuroscience*, 13(2):400-404.
- Clark, H. (1973). The language-as-fixed-effect fallacy. *Journal of Verbal Learning and Verbal Behavior*, 12:335-359.
- Fauconnier, G. (1997). *Mappings in Thought and Language*. Cambridge University Press, Cambridge.
- Fillmore, C. (1982). Frame semantics. In of Korea, L. S., editor, *Linguistics in the Morning Calm*, pages 111-138. Hanshin, Seoul.
- Gallese, V., Fadiga, L., Fogassi, L., and Rizzolatti, G. (1996). Action recognition in premotor cortex. *Brain*, 119:593-609.
- Glenberg, A. M. and Robertson, D. A. (2000). Symbol grounding and meaning: A comparison of high-dimensional and embodied theories of meaning. *Journal of Memory and Language*, 43:379-401.
- Kaschak, M. P. and Glenberg, A. M. (2000). Constructing meaning: The role of affordances and grammatical constructions in sentence comprehension. *Journal of Memory and Language*, 43:508-529.
- Lakoff, G. (1987). *Women, Fire, and Dangerous Things*. University of Chicago Press, Chicago, Illinois.
- Landauer, T. K., Foltz, P. W., and Laham, D. (1998). Introduction to latent semantic analysis. *Discourse Processes*, 25:259-284.
- MacWhinney, B. (1999). The emergence of language from embodiment. In MacWhinney, B., editor, *The emergence of language*, chapter 8. Lawrence Erlbaum Associates, Mahwah, New Jersey.
- Nyberg, L., Habib, R., Macintosh, A., and Tulving, E. (2000). Reactivation of encoding-related brain activity during memory retrieval. In *Proceedings of the National Academy of Sciences*, volume 97, pages 11120-11124, USA.
- Nyberg, L., Petersson, K.-M., Nilsson, L.-G., Sandblom, J. J. J., Aberg, C., and Ingvar, M. (2001). Reactivation of motor brain areas during explicit memory for actions. *Neuroimage*, 14:521-528.
- Porro, C. A., Francescato, M. P., Cettolo, V., Diamond, M. E., Baraldi, P., Zuian, C., Bazzocchi, M., and di Prampero, P. E. (1996). Primary motor and sensory cortex activation during motor performance and motor imagery: a functional magnetic resonance imaging study. *Journal of Neuroscience*, 16:7688-7698.
- Pulvermüller, Friedman, Haerle, M., and Hummel, F. (2001). Walking or talking?: Behavioral and neurophysiological correlates of action verb processing. *Brain and Language*, 78:143-168.
- Rizzolatti, G., Fadiga, L., Gallese, V., and Fogassi, L. (1996). Premotor cortex and the recognition of motor actions. *Cognitive Brain Research*, 3:131-141.
- Tettamanti, M., Buccino, G., Saccuman, M. C., Gallese, V., Danna, M., Perani, D., Cappa, S. F., Fazio, F., and Rizzolatti, G. Sentences describing actions activate visuomotor execution and observation systems.

Embodied Verbal Semantics: lexical matching task

Wheeler, M. E., Petersen, S. E., and Buckner, R. L. (2000). Memory's echo: Vivid remembering reactivates sensory-specific cortex. In *Proceedings of the National Academy of Sciences*, volume 97, pages 11125-11129, USA.