Reconciling Embodied and Distributional Accounts of Meaning in Language

Mark Andrews
Division of Psychology
Nottingham Trent University

Stefan Frank and Gabriella Vigliocco
Department of Cognitive, Percetual and Brain Sciences
University College London

Abstract

Over the past 15 years, there have been two increasingly popular approaches to the study of meaning in cognitive science. One, based on theories of embodied cognition, treats meaning as a simulation of perceptual and motor states. An alternative approach treats meaning as a consequence of the statistical distribution of words across spoken and written language. On the surface, these appear to be opposing scientific paradigms. In this review, we aim to show how recent cross-disciplinary developments have done much to reconcile these two approaches. The foundation to these developments has been the recognition that intra-linguistic distributional and sensory-motor data are interdependent. We describe recent work in philosophy, psychology, cognitive neuroscience and computational modeling that are all based on or consistent with this conclusion. We conclude by considering some possible directions for future research that arise as a consequence of these developments.

1 Introduction

How meaning in language is represented and learned is a foundational problem both in cognitive science and in language engineering research. In cognitive science particularly, a major recent development in the theoretical treatment of meaning in language has been based on theories of embodied cognition. At its most general, this account holds that meaning in language is based on perceptual and motor states that arise from our direct sensory experience and actions in the world. As we will describe, this perspective has lead to considerable empirical progress in cognitive psychology and cognitive neuroscience. In parallel with this development, however, there has also been considerable progress in what are termed distributional models of meaning: computational models that derive semantic representations solely on the basis of statistical distributions over spoken and written language. In distributional models, knowledge of meaning is knowledge of statistical patterns. As such, meaning is literally disembodied. Words are symbols whose statistical behaviors may or may not be related to those of other symbols, but words do not relate to any perceptual or action states of the body.

On the surface, there appears to be a troubling divide or a disunity between the embodied and distributional approaches to meaning. It is particularly troubling because the schism is largely along disciplinary lines. Embodied cognition has become increasingly popular in the experimental disciplines of psychology and cognitive neuroscience, while distributional models have become likewise popular in computational disciplines such as cognitive modelling, computational linguistics and language engineering. This bodes ill for inter-disciplinary approaches to the study of embodiment and language. Indeed, the theme of this special issue of *TopiCS* is precisely to promote such approaches, and particularly to bridge the evident gap between the experimental and computational work in this area.

In this review, we aim to describe a range of recent cross-disciplinary developments that we believe do much to reconcile the differences between the embodied and the distributional approaches to meaning in language. These developments are genuinely cross-disciplinary, having emerged independently in different fields such as philosophy, experimental psychology, cognitive neuroscience, computational modeling and computational linguistics. While the divide between embodied and distributional accounts has not been resolved in its entirety, the recent developments that we describe have done much to minimize the disunity between the two approaches. We conclude by outlining some of the directions of future research of this kind that we believe will be most promising.

2 A Schism in the Study of Meaning

2.1 Embodiment and Meaning

The embodied view of cognition (e.g., Barsalou, 1999, 2008), when applied to the problem of language understanding, holds that meaning is represented in terms of simulations of the perceptual and motor states of a goal-directed agent in an environment. For the problem of word meanings in particular, the embodied cognition account is that the very perceptual states that arise from the perception of an object, for example, and the very motor states involved with action toward or interaction with that object, are simulated whenever we use the word that refers to that object. This account also explicitly holds that the actual neural states that underlie these perceptual and motor states are themselves re-simulated whenever we use this word.

Considerable empirical evidence to support this view has been obtained (for comprehensive reviews, see e.g., Barsalou, 2008; Pecher & Zwaan, 2005). For example, Barsalou, Solomon, and Wu (1999) asked subjects to list the properties of nouns like *watermelon* versus noun-phrases like *half watermelon*. In the latter condition, but not the former, internal properties like *red*, *seeds* and *sweet* were often listed, suggesting that the noun-phrase representation is simulation based. In Pecher, Zeelenberg, and Barsalou (2003), modality switching costs in property verification tasks were observed: Subjects were slower in verifying an auditory property of an object (e.g. blender-loud) when they had just verified a gustatory one (e.g. cranberries-tart) than another auditory property (e.g. leaves-rustling). Meteyard, Bahrami, and Vigliocco (2007) showed that hearing motion-related verbs interferes with the processing of visual motion. Meteyard, Zokaei, Bahrami, and Vigliocco (2008) found the reverse effect: visual motion interfering with the processing of motion-related verbs.

In a comprehensive overview of embodiment research from cognitive neuroscience, Martin (2007) showed that when subjects conceptualize the salient perceptual and motor properties of words, the same sensory and motor systems are activated as when the referents of these words are being perceived. Comparable results were found for action words and their representation in motor cortices (Hauk, Johnsrude, & Pulvermüller, 2004; Tettamanti et al., 2005; Aziz-Zadeh, Wilson, Rizzolatti, & Iacoboni, 2006). Likewise, Goldberg, Perfetti, and Schneider (2006); Simmons et al. (2007); Kan, Barsalou, Solomon, Minor, and Thompson-Schill (2003) showed that the same somatosensory, motor, and premotor cortical regions were active for the verification of the properties of concepts as when the objects themselves are being perceived or acted upon.

2.2 Distributional Statistics and Meaning

In direct contrast to the embodied cognition approach, distributional accounts of meanings seem to eschew all mention of either the body or perception. According to these accounts, the meaning of words is learned from their statistical distribution across language. As Firth (1957), one of the early proponents of this view, famously summarized "You shall know a word by the company it keeps" (p.11). By this statement, word meanings seem to arise solely from intra-linguistic or word-to-word relationships. The embodied conception of the grounding of word meanings in our perception of the world appears to have no place in this approach.

Just as research in embodied approaches to meaning has grown considerably within the last 10 to 15 years, so too have distributional accounts of meaning come to play an increasingly dominant role, especially in computational models of semantic representation. Two of the earliest such models in cognitive science were Hyperspace Analog of Language (HAL) (e.g., Burgess & Lund, 1997) and Latent Semantic Analysis (LSA) (e.g., Landauer & Dumais, 1997). In HAL, the statistical distribution of each word in a large corpus was described by its frequency of co-occurrence with other words. Lund and Burgess (1996) showed that on the basis of this information, semantic categories like *animals*, *body parts*, *geographical regions*, etc., are easily distinguishable from one another. Moreover, the distances between word-pairs in the HAL model correlate positively with semantic priming effects on reaction times in a lexical decision task. LSA defined the statistical distribution of words in terms of whether, and with what frequency, each word occurs in each document in a large collection of text documents. As with HAL, semantic categories and meaningful word-pairs were easily identified. LSA also performed similarly to human participants in a synonym recognition test.

In more recent work, Griffiths, Steyvers, and Tenenbaum (2007) applied a set of Bayesian models of distributional statistics to a wide range of semantic phenomena, and were able to model semantic similarity and word-associations to greater extent than LSA. In addition, they modeled the learning and representation of polysemy, ambiguity resolution, text gist extraction, and the learning and representing of a continuum between syntactic and semantic categories. Using an alternative computational approach, Jones and Mewhort (2007) provided a model of the lexicon derived from distributional statistics. With this, they were able to account for empirical data from classic experiments on semantic typicality, categorization, priming, and semantic effects in sentence completions.

Although the link between distributional models and cognitive neuroscience is not well developed as yet, in pioneering fMRI work, Mitchell et al. (2008) were able to predict distributions of voxel-level neural activation on the basis of the distributional statistics of nouns.

3 Reconciling Embodied and Distributional Accounts

As presented, the embodied and distributional accounts appear to be opposing paradigms. Indeed this is precisely how they are portrayed by some. Glenberg and Robertson (2000), for example, argue against the distributional approach and in favor of the embodied approach. For them, appealing to the arguments of Searle (1980) and Harnad (1990), the fatal but inevitable flaw of distributional accounts of meaning is that the symbols are not grounded. Others doubt that distributional approaches have any relevance to theories of cognition beyond that of being methodological tools (e.g., Perfetti, 1998). By contrast, distributional approaches, while certainly identifying themselves as theories of cognition (e.g., Landauer, 1999a), are often either indifferent to embodiment research, or else reject the necessity of embodied cognition for understanding meaning in language (e.g., Landauer, 1999b).

This state of either splendid isolation or conflict is not welcome, not least because the schism between the two approaches is largely along disciplinary and methodological lines. The embodied approach has become increasingly popular in the experimental sub-fields of cognitive science (e.g., experimental psychology, cognitive neuroscience), while the distributional approach has become equally popular in the computational modelling sub-fields (e.g., artificial intelligence, machine learning, computational linguistics). Moreover, the disunity between the two approaches clearly reveals their obvious strengths and weaknesses. Embodied approaches are challenged to adequately account for abstract concepts and the meanings of words for which tangible physical referents do not exist. These problems do not arise with distributional approaches. On the other hand, while distributional approaches may be able to infer that, for example, the terms *pen* and *write* are related, they would be challenged to infer that they refer to everyday objects and actions. Obviously, these issues do not arise with embodied approaches.

The complimentary nature of these strengths and weakness does not seem coincidental, however, and it may be the means to reconcile the two approaches. Their complimentary characters seem to arise from the fact that both approaches are focused on one of two interdependent data-types, either data in the world or data in the discourse of language itself. That these data-types are interdependent allows one to bootstrap from the other. In fact, it is the emerging realization that these two data-types are not independent, but mutually reinforcing, that has been the foundation to the reconciliatory developments to which we now turn [1].

3.1 Philosophy

A central tenet of embodied cognition generally (e.g., Wilson, 2002), is that cognitive agents structure their physical environment so that it acts as an extension of their own cognitive system. Clark (2006) has presented a compelling account of how language can be seen as another (literally) physical environment that agents may perceive and act upon. By this account, amongst other things, language provides a new source of perceptual data and new targets for action, in a manner identical to that of any other modality. It provides "a new realm of perceptible objects ... upon which to target (our) more basic capacities of statistical and associative learning" (p. 371). In other words, language is an extension of the physical environment generally, and one that we may perceive (by language comprehension) and act upon (by language production), just as we do with any physical environment.

There are at least two important consequences of this perspective. First, if language is just another environment, distributional approaches to language do not differ fundamentally from embodied approaches to cognition generally, with the exception of concentrating more exclusively on one source of environmental data. Second, because language is a man-made environment (the *ultimate artifact* according to Clark, 1998), this allows for much greater structuring so that it acts as an extension or complement to our own cognition. For Clark (2006), this allows us to use language to model the world. We use language to *re-present*, to synthesize, to abstract from otherwise bafflingly complex data and relations in the world. Likewise, Dove (2011) argues that our experience with language as an 'amodal symbol system that is built on an embodied substrate' (p. 8) enhances our ability to learn concepts, in particular abstract concepts.

3.2 Psychology

If language is just another kind of physical environment, integrating language-based statistical data with perceptual data generally is a kind of multimodal or sensory-motor integration (e.g., Körding & Wolpert, 2004; Deneve & Pouget, 2004). Recent evidence from experimental psychology supports this view. For example, Louwerse and Jeuniaux (2010) showed that regularities in both linguistic and visual data influence the processing of conceptual information presented either pictorially or verbally, with predictable interactions depending on the precise nature of the task. In particular, conceptual judgements about *both* pictures of objects and pairs of words were influenced by *both* the usual spatial relations between objects in the world (e.g., cars are above roads, bridges are above rivers, etc.) and the usual order of words referring to these objects. This also holds for pictures of attics and basements. The linguistic information was more facilitatory for judgements about words, the spatial information more so for judgements about pictures. Likewise, the spatial information was more facilitatory for judgements about iconic relationships, while the linguistic regularities were more relevant to semantic judgements.

There is also growing evidence in support of Clark's (2006) hypothesis that we use language as a model of, or proxy for, the world. In fact, this is highly related to the *symbol interdependency*

[1] While the studies that we describe are all recent developments, they are not without precedent. For example, the well known dual-coding theory of Paivio (Paivio, 1971, 1986/1990) proposed a distinction between knowledge acquired by way of sensory processes and that acquired through language, viewing both as necessary and complementary of one another. More recently, Deacon (1997) proposed an account of cognition that emphasized how linguistic knowledge is a complex network of abstract symbolic relationships that connects at its periphery to sensory experience. As important as these ideas are, we regard them more as precursors to the current discussion than as substantial parts of it.

hypothesis of Louwerse (2007, 2011; Louwerse & Jeuniaux, 2008) that describes language as encoding relations in the world, including embodied relationships. By this account, we can learn and represent the world by way of learning these *intra*-linguistic relationships. Louwerse (2011) describes how the perceptual relationships underlying the modality-switching task of Pecher et al. (2003) (see Section 2.1) could be learned from distributional relationships in language. Louwerse (2008) similarly shows that spatial relations are encoded in word-order statistics. Riordan and Jones (2011) provide considerable evidence that distributional statistics often provide much of the information provided in sensory-motor feature-based data. Lynott and Connell (2010) propose a model of how distributional linguistic information guides the formation of what are ultimately simulated conceptual combinations.

3.3 Cognitive neuroscience

Recent reviews of the cognitive neuroscience of semantic memory by Meteyard, Cuadrado, Bahrami, and Vigliocco (2012) and by Binder and Desai (2011) show that there is growing evidence against what they describe as *strong embodiment* theories of neural representation of conceptual knowledge, i.e., full analogical simulations of actual sensory-motor experience. In other words, although it may be that under certain conditions speakers do run full simulations[2] of their experience and thereby engage primary sensory and motor system (e.g., Gallese & Lakoff, 2005), this is not always necessary. According to Meteyard et al. and Binder and Desai, existing evidence strongly suggests that multimodal integration in convergence zones of increasing complexity and increasing abstraction is an essential component of semantic representation in the brain. These views, in which information from different modalities is hierarchically integrated, lend themselves naturally to the assumption that linguistically derived information is also integrated to form semantic representations. Likewise, they are also in line with the work of Louwerse (2007, 2011); Louwerse and Jeuniaux (2008) that claims that full embodied simulations are not always necessary for language processing and conceptual reasoning.

3.4 Modeling of word meanings

On the basis of the same general assumptions that motivated the experimental work described in Section 3.2, namely that distributional statistics in language and perceptual data in the world are two complimentary and correlated data types, there have been numerous recent attempts to model the learning of semantic representations using these two data types in parallel. Andrews, Vigliocco, and Vinson (2009) used a Bayesian model based on a generalization of the previously mentioned work of Griffiths et al. (2007) to learn couplings of clusters of sensory-motor features with discourse topics in a large corpus. In so doing, this model revealed the correlations both within and *between* these two data-types. The semantic representations that were thus learned were more similar to human semantic representations as revealed from different empirical datasets. In related work, Steyvers (2010) also showed an improved generalization performance using a Bayesian model that learned from distributional and sensory-motor data in parallel.

One of the attractive features of these models is that they permit inference from one data type to the other. In other words, we may map arbitrary words to sensory-motor features, thereby addressing the issue of symbol-grounding. Louwerse's symbol interdependency hypothesis proposes symbol-to-symbol and symbol-to-world relationships, although ultimately any symbol may be grounded through chains. These models provide computational demonstrations of this phenomenon. This is demonstrated most comprehensively by Johns and Jones (2012). Using a global memory model, they model words as joint distributions over sensory-motor features and co-occurrence patterns with texts. By so doing, and on the basis of the redundancy of perceptual and distributional data, they are able to infer sensory-motor features for arbitrary words. From this, they predict behaviours in a range of data-sets from embodied cognition experiments.

[2] Defined by Meteyard et al. (2012) as *the re-creation of direct experience through the modulation of activity in primary sensory and motor areas.*

The computational models just described all use speaker-generated feature norms (from e.g., Vigliocco, Vinson, Lewis, & Garrett, 2004; McRae, de Sa, & Seidenberg, 1997) as a rough proxy to real sensory-motor data. This is not ideal, and it is questionable to what extent feature norms accurately reflect the real data that we experience. In view of this limitation, and the fact that speaker-generated feature norms are minimal data-sets that limit the scale of computational models, Bruni, Tran, and Baroni (2011) explored the integration of perceptual and distributional data using the actual visual information in images. In so doing, they combine state-of-the-art statistical work from computer vision research with that of computational linguistics and language engineering. As they demonstrate, distributional models are improved by augmenting them with with visual information related to the referents of words. In particular, this allows them to more easily learn the semantic similarities between concrete words.

3.5 Modeling meaning beyond words

Computational models of sentence-level meaning that aim to take world knowledge into account face the problem that a vast amount of knowledge is (potentially) relevant yet not available in a form that easily affords statistical modeling. Such models are therefore necessarily restricted to a limited world domain, which requires only an equally domain-limited language. This approach was used in sentence-comprehension models by St.John and McClelland (1990), Nenov and Dyer (1994), and Frank, Haselager, and Van Rooij (2009), as well as in several demonstrations of language comprehension in the cognitive robotics literature (e.g., Madden, Hoen, & Dominey, 2009). Despite these limitations, models that live in a limited world and are exposed to a miniature language may provide valuable insight into the integration of linguistic and embodied experience. The restricted and artificial nature of the modeled experience, however, makes it challenging to quantitatively compare model predictions to psycholinguistic data. As a potential exception, the model by Chang and Gurevich (2004) learns particular linguistic constructions (such as 'PERSON throw OBJECT') from distributional patterns in small part of the CHILDES database that is hand-annotated with semantic and contextual information. Although its input is more realistic than that of other sentence-level models, the lack of appropriately annotated data severely restricts the range of linguistic expressions the model can deal with.

In general, sentence-level models are confronted with challenges that do not arise at the word level. These include the issues of productivity (a relatively small set of words can be combined to form an unbounded number of possible sentences) and compositionality (the meaning of a sentence depends non-trivially on the meaning of its words and how they are combined). Consequently, modeling linguistic meaning by integrating distributional and embodied approaches is much more challenging at the sentence level than at the word level. Nevertheless, there is no fundamental reason why it would be impossible, or even just infeasible. For example, a model that learns linguistic statistics from large text corpora could be enhanced with domain-limited but realistic world knowledge. Its processing of realistic language input within its knowledge domain could then be directly compared to humans processing the same sentences, allowing for quantitative model evaluation.

4 Conclusion

In this review, we have tried to show how recent developments have done much to reconcile two general approaches to the study of meaning in language. What we have called the embodied approach treats meaning as a simulation of perceptual and motor states. The parallel trend that we have called the distributional approach treats meaning as a consequence of the statistical distribution of words across spoken and written languages. The foundation and common theme to the reconciliatory developments that we have described has been the recognition that intra-linguistic distributional and sensory-motor data are interdependent data-types. We don't claim that these developments have resolved all problems. However, we feel that there is no longer a sense in which the embodied and distributional approaches are opposing paradigms.

Some important issues still remain even if we accept the premise that distributional and sensory-motor data are interdependent. One concerns the exact nature of their inter-dependence. Clark's (2006) argument that language is an environment that we manipulate leaves open the question of whether we should view language as just one more data source from which we learn semantic information, or a special type of data that is a proxy to that of the physical world. The symbol interdependency hypothesis of Louwerse (2007, 2011); Louwerse and Jeuniaux (2008) and the emphasis on convergence zones in recent cognitive neuroscience is perhaps ambivalent between these two possibilities. Another, more practical, issue is that currently, sensory-motor data does not lend itself to computational modeling to the same extent as does language data. Massive text corpora are easily available and these are, relatively speaking, simple to use as the data in statistical models. This ease of use, rather than any theoretical commitment, was precisely the reason for the increased use of distributional models that began in the 1990s. By contrast, multimodal sensory-motor data is either not available at all, or presents formidable preprocessing challenges before it can be used in language modeling. Advances in engineering, and especially in robotics, seem necessary for the development of this kind of modeling. In this respect, progress in the field of robotics may be both the cause and consequence of a full convergence of the embodied and distributional approaches to meaning.

References

Andrews, M., Vigliocco, G., & Vinson, D. P. (2009). Integrating experiential and distributional data to learn semantic representations. *Psychological Review*, *116*(3), 463-498.

Aziz-Zadeh, L., Wilson, S. M., Rizzolatti, G., & Iacoboni, M. (2006). Congruent embodied representations for visually presented actions and linguistic phrases describing actions. *Current Biology*, *16*, 1818-1823.

Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences*, *22*, 577-660.

Barsalou, L. W. (2008). Grounded cognition. *Annual Review of Psychology*, *59*, 617-645.

Barsalou, L. W., Solomon, K., & Wu, L. (1999). Perceptual simulation in conceptual tasks. In *Cultural, Typological, and Psychological Perspectives in Cognitive Linguistics: The Proceedings of the 4th Conference of the International Cognitive Linguistics Association* (Vol. 3, p. 209-228).

Binder, J. R., & Desai, R. H. (2011). The neurobiology of semantic memory. *Trends In Cognitive sciences*, *15*(11), 527-536.

Bruni, E., Tran, G. B., & Baroni, M. (2011). Distributional semantics from text and images. In *Proceedings of the EMNLP 2011 Geometrical Models for Natural Language Semantics (GEMS 2011) Workshop* (p. 22-32). Stroudsburg, PA: Association for Computational Linguistics.

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

Chang, N., & Gurevich, O. (2004). Context-driven construction learning. In K. D. Forbus, D. Gentner, & T. Regier (Eds.), *Proceedings of the 26th Annual Meeting of the Cognitive Science Society* (pp. 204–209). Hillsdale, NJ: Erlbaum.

Clark, A. (1998). *Being there: Putting brain, body, and world together again*. Cambridge, MA: MIT Press.

Clark, A. (2006). Language, embodiment, and the cognitive niche. *Trends in Cognitive Sciences*, *10*(8), 370-374.

Deacon, T. W. (1997). *The Symbolic Species: the Co-Evolution of Language and the Brain*. New York, NY: W. W. Norton.

Deneve, S., & Pouget, A. (2004). Bayesian multisensory integration and cross-modal spatial links. *Journal of Physiology (Paris)*, *98*, 249-258.

Dove, G. (2011). On the need for embodied and dis-embodied cognition. *Frontiers in Psychology*, *1:242*.

Firth, J. R. (1957). A synopsis of linguistic theory 1930-1955. In *Studies in Linguistic Analysis (special volume of the Philological Society, Oxford)* (p. 1-32). Oxford: Blackwell.

Frank, S. L., Haselager, W. F. M., & Van Rooij, I. (2009). Connectionist semantic systematicity. *Cognition*, *110*, 358–379.
Gallese, V., & Lakoff, G. (2005). The brain's concepts: The role of the sensory-motor system in conceptual knowledge. *Cognitive Neuropsychology*, *22*(3-4), 455-479.
Glenberg, A. M., & Robertson, D. A. (2000). Symbol grounding and meaning: A comparison of high-dimensional and embodied theories of meaning. *Journal of Memory and Language*, *43*(3), 379-401.
Goldberg, R., Perfetti, C., & Schneider, W. (2006). Perceptual knowledge retrieval activates sensory brain regions. *Journal Of Neuroscience*, *26*(18), 4917-4921.
Griffiths, T. L., Steyvers, M., & Tenenbaum, J. B. (2007). Topics in semantic representation. *Psychological Review*, *114*(2), 211-244.
Harnad, S. (1990). The symbol grounding problem. *Physica D*, *42*, 335-346.
Hauk, O., Johnsrude, I., & Pulvermüller, F. (2004). Somatotopic representation of action words in human motor and premotor cortex. *Neuron*, *22*, 301-307.
Johns, B. T., & Jones, M. N. (2012). Perceptual inference through global lexical similarity. *Topics in Cognitive Science*, *4*, 103-120.
Jones, M. N., & Mewhort, D. J. K. (2007). Representing word meaning and order information in a composite holographic lexicon. *Psychological Review*, *114*(1), 1-37.
Kan, I., Barsalou, L., Solomon, K., Minor, J., & Thompson-Schill, S. (2003). Role of mental imagery in a property verification task: fMRI evidence for perceptual representations of conceptual knowledge. *Cognitive Neuropsychology*, *20*(3-6), 525-540.
Körding, K., & Wolpert, D. (2004). Bayesian integration in sensorimotor learning. *Nature*, *427*, 244-247.
Landauer, T. (1999a). Latent semantic analysis: A theory of the psychology of language and mind. *Discourse Processes*, *27*(3), 303-310.
Landauer, T. (1999b). Latent Semantic Analysis (LSA), a disembodied learning machine, acquires human word meaning vicariously from language alone. *Behavioral and Brain Sciences*, *22*(4), 624-625.
Landauer, T., & Dumais, S. (1997). A solutions to Plato's problem: The Latent Semantic Analyis theory of acquistion, induction and representation of knowledge. *Psychological Review*, *104*, 211-240.
Louwerse, M. M. (2007). Symbolic or embodied representations: A case for symbol interdependency. In T. Landauer, D. McNamara, S. Dennis, & W. Kintsch (Eds.), *Handbook of Latent Semantic Analysis* (p. 107-120). Mahwah, NJ: Erlbaum: Erlbaum.
Louwerse, M. M. (2008). Embodied relations are encoded in language. *Psychonomic Bulletin and Review*, *15*(4), 838-844.
Louwerse, M. M. (2011). Symbol Interdependency in Symbolic and Embodied Cognition. *Topics In Cognitive Science*, *3*(2), 273-302.
Louwerse, M. M., & Jeuniaux, P. (2008). Language comprehension is both embodied and symbolic. In A. C. G. M. de Vega A. Glenberg (Ed.), *Symbols and embodiment: Debates on meaning and cognition* (p. 309-326). Oxford, UK: Oxford University Press.
Louwerse, M. M., & Jeuniaux, P. (2010). The linguistic and embodied nature of conceptual processing. *Cognition*, *114*(1), 96-104.
Lund, K., & Burgess, C. (1996). Producing high-dimensional semantic spaces from lexical co-occurrence. *Behavior Research Methods, Instrumentation, and Computers*, *28*, 203-208.
Lynott, D., & Connell, L. (2010). Embodied conceptual combination. *Frontiers in Cognition*, *1*(212), 1-14.
Madden, C. J., Hoen, M., & Dominey, P. F. (2009). A cognitive neuroscience perspective on embodied language for human-robot cooperation. *Brain & Language*, *112*, 180–188.
Martin, A. (2007). The representation of object concepts in the brain. *Annual Review of Psychology*, *58*, 25-45.
McRae, K., de Sa, V., & Seidenberg, M. (1997). On the nature and scope of featural representation of word meaning. *Journal of Experimental Psychology: General*, *126*, 99-130.

Meteyard, L., Bahrami, B., & Vigliocco, G. (2007). Motion detection and motion verbs - Language affects low-level visual perception. *Psychological Science*, *18*(11), 1007-1013.
Meteyard, L., Cuadrado, S. R., Bahrami, B., & Vigliocco, G. (2012). Coming of Age: A review of embodiment and the neuroscience of semantics. *Cortex*, *48*(7), 788-804.
Meteyard, L., Zokaei, N., Bahrami, B., & Vigliocco, G. (2008). Visual motion interferes with lexical decision on motion words. *Current Biology*, *18*(17), R732-R733.
Mitchell, T. M., Shinkareva, S. V., Carlson, A., Chang, K.-M., Malave, V. L., Mason, R. A., et al. (2008). Predicting human brain activity associated with the meanings of nouns. *Science*, *320*(5880), 1191-1195.
Nenov, V. I., & Dyer, M. G. (1994). Perceptually grounded language learning: Part 2 — DETE: a neural/procedural model. *Connection Science*, *6*, 3–128.
Paivio, A. (1971). *Imagery and verbal processes.* New York: Holt, Rinehart, and Winston.
Paivio, A. (1990). *Mental representations: A dual-coding approach.* New York: Oxford University Press. (Original work published 1986)
Pecher, D., Zeelenberg, R., & Barsalou, L. (2003). Verifying different-modality properties for concepts produces switching costs. *Psychological Science*, *14*(2), 119-124.
Pecher, D., & Zwaan, R. A. (Eds.). (2005). *Grounding cognition: The role of perception and action in memory, language, and thinking*. Cambridge,UK: Cambridge University Press.
Perfetti, C. (1998). The limits of co-occurrence: Tools and theories in language research. *Discourse Processes*, *25*(2-3), 363-377.
Riordan, B., & Jones, M. N. (2011). Redundancy in Perceptual and Linguistic Experience: Comparing Feature-Based and Distributional Models of Semantic Representation. *Topics in Cognitive Science*, *3*(2), 303-345.
Searle, J. R. (1980). Minds, brains and programs. *Behavioral and Brain Sciences*, *3*(3), 417-457.
Simmons, W. K., Ramjee, V., Beauchamp, M. S., McRae, K., Martin, A., & Barsalou, L. W. (2007). A common neural substrate for perceiving and knowing about color. *Neuropsychologia*, *45*(12), 2802-2810.
Steyvers, M. (2010). Combining feature norms and text data with topic models. *Acta Psychologica*, *133*, 234-243.
St.John, M. F., & McClelland, J. L. (1990). Learning and applying contextual constraints in sentence comprehension. *Artificial Intelligence*, *46*, 217–257.
Tettamanti, M., Buccino, G., Saccuman, M. C., Gallese, V., Danna, M., Scifo, P., et al. (2005). Listening to action-related sentences activates fronto-parietal motor circuits. *Journal of Cognitive Neuroscience*, *17*(2), 273-281.
Vigliocco, G., Vinson, D. P., Lewis, W., & Garrett, M. F. (2004). Representing the meanings of object and action words: The featural and unitary semantic space hypothesis. *Cognitive Psychology*, *48*, 422-488.
Wilson, M. (2002). Six views of embodied cognition. *Psychonomic Bulletin & Review*, *9*(4), 625-636.