

Toward a theory of semantic representation

GABRIELLA VIGLIOCCO, LOTTE METEYARD, MARK ANDREWS, AND
STAVROULA KOUSTA*

University College London

Abstract

We present an account of semantic representation that focuses on distinct types of information from which word meanings can be learned. In particular, we argue that there are at least two major types of information from which we learn word meanings. The first is what we call experiential information. This is data derived both from our sensory-motor interactions with the outside world, as well as from our experience of own inner states, particularly our emotions. The second type of information is language-based. In particular, it is derived from the general linguistic context in which words appear. The paper spells out this proposal, summarizes research supporting this view and presents new predictions emerging from this framework.

Keywords

word meaning, semantic representation, conceptual structure, embodiment, Bayesian modelling

1. Introduction

The ability to use language to refer to the here and now, the past, the future, the hypothetical and the imaginary is at the core of all human endeavours and is a uniquely human faculty. Underpinning this ability is the ability to develop mental representations for things, events, properties

* Correspondence address: Gabriella Vigliocco, Cognitive, Perceptual and Brain Sciences Research Department, Division of Psychology and Language Sciences, University College London, 26 Bedford Way, London, WC1H 0AP, UK. E-mail: g.vigliocco@ucl.ac.uk. The work reported here has been supported by an Economic and Social Research Council Studentship to L. M. by the European Union (FP6-2004-NESTPATH) Grant 028714, and by the UK Biotechnology and Biological Sciences Research Council (BBSRC) Grant 31/S18 to G. V.

and abstract notions and to establish connections between these representations and the specific linguistic forms used in one's language. This paper discusses our proposal of how meaning is represented in the mind/brain across domains of knowledge, both concrete and abstract. It is structured in the following manner. We first provide an overview of our theory, spelling out our assumptions and highlighting elements of convergence with other approaches as well as elements of novelty. We then present a review of the evidence (from cognitive psychology, cognitive neuroscience and computational modelling) in favour of our view.

2. The joint contribution of experience and language in shaping meaning

In continuity with our previous work (Vigliocco et al. 2004) and with a large body of literature, one core assumption of our proposal is an embodied view of cognition. In contrast to amodal theories of semantic and conceptual representation, e.g. Jackendoff (2002), embodied theories of cognition propose that cognition is grounded in bodily states, modal simulations, and situated action (Barsalou 1999; Barsalou et al. 2003; Decety and Grezes 2006; Glenberg and Robertson 2000; Gibbs 2006; Rizzolati and Craighero 2004; Zwaan 2004). Despite some disagreements about the exact nature of the link between sensory-motor experience and semantic information (see, for example, Gallese and Lakoff 2005 vs. Vigliocco et al. 2004), all embodied theories share the core assumption that the representation and processing of semantic information automatically recruits, in some form or other, the same neural systems that are engaged during perception and action. Recent work, some of which is described in more details below, has provided evidence for such a link between semantic and sensory-motor information, either by showing that perception/action affects semantic processing (Kaschak et al. 2005; Kaschak et al. 2006; Meteyard et al. 2008) or that semantic processing affects perception/action (Meteyard et al. 2007; Glenberg and Kaschak 2003).

Our proposal departs however from the majority of current theories because of the following reasons. First, it attempts to provide an account of how semantic knowledge is represented across all domains of knowledge. Second, we emphasize the role of affective, or emotional, information as another type of experiential information that is foundational (i.e. primary and necessary) in learning and representing meanings, especially for abstract words. Third, our proposal sharply contrasts with other embodiment views that consider semantic representation as uniquely arising from experiential information in that we hypothesize that language itself also provides vital information from which to learn semantic representation across all domains of knowledge. The emphasis on linguistic and

especially affective information is motivated on the grounds that sensory-motor information is insufficient in accounting for learning and representation of especially abstract words. Finally, we provide for the first time an explicit account of how experiential and linguistic information are combined in semantic representation.

While embodied approaches emphasizing sensory-motor information can be straightforwardly applied to the representation and processing of concrete word meanings, it is less obvious how an embodied account can deal with abstract word meanings, which have traditionally been considered to be within the purview of purely verbal systems (e.g. Paivio 1986). In one approach, which originates in work in cognitive linguistics, abstract concepts are grounded metaphorically in embodied and situated knowledge (Lakoff and Johnson 1980, 1999; Gibbs 1994; see also Coulson 2000; Turner and Fauconnier 2000). The proposal is that abstract knowledge is viewed as originating in “conceptual metaphors” (i.e. the use of a concrete conceptual domain of knowledge to describe an abstract conceptual domain). For example English consistently uses language concerning *throwing* and *catching* to describe communication of ideas. In this view, representation of abstract concepts in the mind/brain is grounded in the representation of concrete knowledge, which in turn is grounded in our sensory and motor experience of the world. Although there is increasing evidence that metaphors play a role in the conceptualization of some abstract domains (Boroditsky and Ramscar 2002; Gibbs 2006), it is a matter of controversy whether they are foundational in the development (and subsequent representation) of abstract concepts and word meanings or whether they provide structure to pre-existing conceptual content (Barsalou 1999; but see Glenberg et al. 2008). Moreover, even if unequivocal evidence in favor of metaphorical grounding were to be found, it remains an open issue how far such a mechanism could go in accounting for the variety of abstract knowledge humans master and how the system decides among multiple metaphors that can be applied to the same domain (e.g. time can be conceived like a spatial domain, however, other metaphors are also possible, e.g. ‘time is a limited resource’, as in “time is running out” “we don’t have any time left”). In the proposal we spell out here, although we do leave open the possibility that some abstract domains may be grounded in perception and action via conceptual metaphors, we further suggest two additional types of information that may play a special role in the learning and representation of abstract words, namely, affective and linguistic information.

We are proposing that emotion (another type of experiential information) may play a crucial role in the representation and processing of abstract concepts for two main reasons. First, words referring to emotions

or with emotional content are abstract words (see also Altarriba and Bauer 2004 for a related proposal). Second, just like sensory-motor development, emotional development precedes the development of language in children (Bloom 1998). Words that denote emotional states, moods or feelings provide perhaps a crucial example of how a word may refer to an entity that is not observable but resides within the organism. In this manner, the acquisition of words denoting emotions, moods or feelings may actually be a crucial stepping stone in the development of abstract semantic representations. According to Gleitman and colleagues (Gleitman et al. 2005), early word learning is achieved by means of word-to-world mappings (i.e. by observing the situational contingencies of word usage), which is the case for a limited set of words that refer to concrete, basic level concepts. Here we propose that abstract words denoting emotional states, moods or feelings also fall in the same category of words for which a mapping from the word to the world (albeit internal) is possible. Consistent with this hypothesis, words denoting emotional states emerge early in language development, at around 20 months of age, and their rate of acquisition increases rapidly in the third year of life (Bretherton and Beeghly 1982; Wellman et al. 1995). For instance, Ridgeway et al. (1985) report that 76.7% of children aged 18–23 months have acquired the meaning of the words *good*, and *happy*, with approximately half of the children in the same age group have acquired the meaning of the words *tired*, *sad*, *afraid*, and *busy*.

In addition to assuming a critical role for emotion, we further hypothesise that also non-experiential information, namely linguistic information, plays an essential part in learning word meanings. Following a long-standing tradition in computational cognitive science, we are assuming here that important aspects of meaning representation are learnt on the basis of the statistical distribution of words across texts (Schutze 1992; Landauer and Dumais 1997; Burgess and Lund 1997; Lund and Burgess 1996; Griffiths et al. 2007). In our hypothesis, the meaning information that can be implicitly extracted from co-occurrence in text is combined with the *qualitatively different* information derived from sensory-motor experience (Andrews et al. 2009). It is worth mentioning here that, although to-date we have limited our effort to modeling distributional information in unordered context, we believe that word-order and hierarchical syntactic information also provides important constraints to word meaning, as demonstrated in the work of Gleitman and colleagues (e.g. Gleitman et al. 2005). The integration of linguistic and experiential information, as we discuss below, allows us to hook up to the world the linguistic information that would be otherwise ungrounded (Harnad 1990). Whereas the general idea of word meanings being linked to both

experiential and linguistic information has been put forward in one way or another by others (see e.g. Paivio's dual-coding view and more recently the LASS theory developed by Barsalou and colleagues) what is new in our approach is that we, as we will see below, provide a clear hypothesis concerning how experiential and linguistic information can be achieved by humans.

Thus to summarize, the main assumptions we make are as follows:

1. Two classes of information contribute to the representation of both concrete and abstract words: experiential (sensory, motor, but also affective) and linguistic (verbal associations arising through co-occurrence patterns and syntactic information).
2. Differences between concrete and abstract word meanings (as well as within concrete and within abstract word meanings) arise as a result of the proportion and exact type of experiential and linguistic information from which they are derived.
3. The apparent dichotomy between concrete and abstract word meanings arises because of a statistical preponderance for sensory-motor information to underlie concrete word meanings and a preponderance for affective and linguistic information to underlie abstract word meanings. While sensory-motor information is statistically more preponderant for concrete word meanings, affective and linguistic information is statistically more important for abstract word meanings, both for their acquisition and their subsequent representation in the adult system.
4. Humans integrate these different types of information available to them in the physical and linguistic environment in learning semantic representations. This statistical integration, that can be modelled using Bayesian learning and inference, provides a powerful mechanism for learning across all domains of knowledge.

It is important to note here, that just as in our previous work (as in most work from an embodiment perspective), our emphasis is on the *content* of semantic representation, rather than on the *organization* of the semantic space. More specifically, our working hypothesis is that the organization of the semantic system *emerges* from the specific statistical properties of representations: differences in organization come about as consequence of statistical differences in the content of different concepts. Theories based on content of semantic representation (rather than organization) are more naturally linked to neuroanatomical models and hence to the development and testing of predictions concerning what neural representations and networks are shared by language and perception, action and affect, as well as what kind of representations and networks may be

common for words from different domains. Theories concerned with the organization of semantic knowledge, such as holistic network models of semantics (Fodor 1976; Fodor et al. 1980; Levelt 1989; Levelt et al. 1999) and theories discussing differences across domains of knowledge (objects vs. actions or concrete vs. abstract) in terms of characteristics of the semantic space (e.g. hierarchical vs. associative) (Crutch and Warrington 2005; Graesser et al. 1987; Huttenlocher and Lui 1979), instead, have greater difficulty linking with clear neuroanatomical hypotheses as the anatomical and neurophysiological correlates of structure are harder to identify and more open to interpretation. In contrast, the areas of the brain that process particular modalities of information are relatively well established in the neuroscientific literature. In the following section we present evidence for the assumptions above.

3. Evidence for the role of experiential information

3.1. Sensory-motor information

Meteyard and Vigliocco (2008) provide a review of evidence compatible with the foundational role of sensory-motor information in processing language referring to action and perception. To briefly summarise, there is a growing number of behavioural studies showing interactions between sensory-motor systems and language processing. These studies have typically used paradigms that assess the impact of language processing on sensory-motor tasks, and vice versa (e.g. Boulenger et al. 2006; Meteyard et al. 2008). Within cognitive neuroscience, a growing number of imaging and Transcranial Magnetic Stimulation (TMS) studies have provided evidence that the primary motor cortex, and to a substantially lesser extent, primary sensory cortices are engaged in processing words and sentences referring to motion or sensory information (e.g. Buccino et al. 2001; Hauk et al. 2004; Kemmerer et al. 2007; Tettamanti et al. 2005; Vigliocco et al. 2006). One common issue in interpreting results from these studies is the extent to which the reported interactions reflect automatic engagement of shared sensory-motor representations between language comprehension and the sensory-motor tasks. This is because most of the results from behavioural studies are amenable to alternative explanations in terms of attentional/decision biases rather than shared processes/representations. Likewise, much of the imaging literature is amenable to alternative interpretations in terms of ad-hoc recruitment of visual-motor imagery during comprehension. These same criticisms do not easily apply to studies using TMS, which have allowed investigators to explore the time-locked activity of the motor cortex following the comprehension or pro-

duction of action related language. For example, Buccino et al. (2005) found that the amplitude of motor evoked potentials in hand muscles were significantly smaller when listening to hand action sentences, and smaller for the foot muscles when listening to foot action sentences (see also Pulvemüller et al. 2005; Oliveri et al. 2004). This evidence supports the fast and automatic modulation of primary motor areas by words or sentences that describe motor actions.

Using a novel integration of psycholinguistic and psychophysical methods, our group recently provided evidence that comprehending words referring to upward and downwards motion (e.g., *rising*, *dropping*) impacts *low-level* motion perception processes, and vice versa (Meteyard et al. 2007; Meteyard et al. 2008), thus providing evidence for shared representations between the two systems which are automatically engaged in language processing. In particular, by using signal detection theory and visual stimuli with carefully controlled perceptual properties, we were able to disentangle perceptual effects from effects due to decision and executive processes.

In Meteyard et al. (2007), we established that listening to blocks of words referring to motion (up or down) affected motion perception at threshold for random dot kinematograms (RDKs), (Green 1961; Levinson and Sekuler 1976). The use of RDKs is motivated on the grounds that they activate a set of well characterized mechanisms in the early stages of visual processing (e.g. Britten et al. 1992; Nakamura et al. 2003; Scase et al. 1996). Moreover, evidence indicates that activation of MT/V5 is directly modulated by the degree of motion coherence in the display (e.g. Nakamura et al. 2003; Tsushima et al. 2006). Showing that motion *meaning* affects motion detection at threshold for these stimuli provides evidence that processing words referring to motion and processing motion in the display engage area MT/V5. In the experiment, after having established thresholds for each participant, we asked them to carry out the motion detection task (with coherence values for the RDKs set at individual threshold levels) while listening to one of three types of blocks of words: matching (where the word's motion was in the same direction as the motion in the display) mismatching (where the word's motion was in the opposite direction to the display) or control (where the words did not have any upward/downward directionality). Although no effect was present in the analysis of reaction times, results of accuracy data were clear-cut. For perceptual sensitivity (d' , a measure reflecting the ability to separate signal from noise), we found interference (lower d') in the mismatch condition (mismatch between implied word meaning and direction of motion in the RDK), indicating that a mismatch impaired their ability to separate signal from noise. In contrast, for criterion (C , reflecting

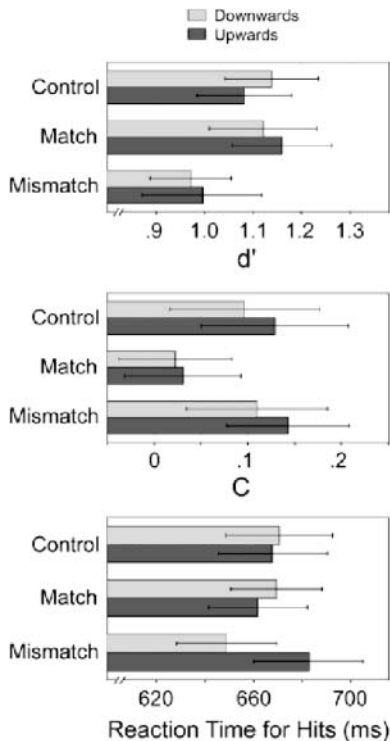


Figure 1. Results of Meteyard, Bahrami, and Vigliocco (2007), permission for reproduction requested. (a) Perceptual sensitivity (d') as a function of match/mismatch between word meaning and motion display. (b) Criterion (C). (c) overall RTs

biases in response decision) we found that participants were more likely to answer “yes” when motion in the display and implied direction of the motion words matched than in the other conditions. These results support a close relationship between the motion representations used in the two tasks. Results from this study are reported in Figure 1.

This study established that the comprehension of words referring to motion engages some of the same processes and representations engaged by motion detection. This conclusion is strengthened by additional findings showing that motion perception affects a linguistic task on words referring to motion. Meteyard et al. (2008) showed that irrelevant motion presented in the background slowed the recognition of words referring to motion when the irrelevant motion was at the threshold of perception, but not suprathreshold (clearly visible). Speakers were asked to carry out a lexical decision task, while ignoring a visual stimulus (RDK) presented

in the background. For one group of subjects the level of motion coherence of the RDK pattern was set at their individual threshold, for other groups, it was set at suprathreshold levels (30, 60 or 90% coherent motion). We found an effect of motion display on RTs, such that subjects were slower at the lexical decision task when the motion implied by the word and the direction of coherent motion in the RDK were mismatching. However, this effect was only present when the coherence level was at threshold, not when it was suprathreshold. For the suprathreshold condition, we instead found a tendency for errors to be more common for motion words in general regardless of match or mismatch between word meaning and display.

These somewhat surprising findings can be interpreted given recent data reported by Tsushima et al. (2006). They showed that irrelevant motion at subliminal and threshold levels interfered with performance in an unrelated visual task (identifying a target letter during rapid serial visual presentation in foveal vision), whereas irrelevant motion at suprathreshold levels did not cause any interference. fMRI data showed that area MT/V5 was more strongly activated when motion was at threshold than when it was suprathreshold. For suprathreshold, motion activation in area MT/V5 was modulated by concomitant activation of dorsolateral prefrontal cortex (DLPFC). The authors interpreted these results to indicate an inhibitory role of DLPFC over area MT/V5. This inhibition, however, would take place only when the signal is clearly perceived and therefore able to initiate suppression, not otherwise. The results of Meteyard et al. (2008) can be interpreted in a similar way. Representations in area MT/V5 are strongly activated by threshold motion signals and this activation interferes with the lexical decision task because the motion processing areas are automatically engaged when reading motion words. Importantly we showed that this interference is not generalized to all motion words, regardless of direction, but is specific to conditions in which direction of the display mismatches with direction of motion implied by the words and the two signals conflict suggesting a high degree of specificity in the neural populations engaged by motion meaning. Suprathreshold motion would be, instead, inhibited by executive functions. There is a cost to this process of inhibition, as a non-selective suppression of activation in MT/V5 causes greater difficulty (more errors) for all motion words.

Taken together, our results cannot be easily accounted for in terms of executive or attentional processes thus strongly support the view that language comprehension does automatically activate the low-level visual processing areas activated during motion perception, as argued by embodiment theories. However, our data and most of the other evidence

available are silent with respect whether activation of sensory-motor representations is *necessary* for comprehension to occur. We have discussed this crucial question elsewhere (Meteyard and Vigliocco 2008) recognising that, on the basis of the available literature, activation of shared representations between perception/action and language processing for adult speakers may not be necessary. The situation, however, may well be different during development (see Glenberg and Gallese, submitted, for evidence and arguments in favor of this latter view).

3.2. *Affective information*

In the cognitive psychology and cognitive neuroscience literature, a clear divide is traditionally drawn between emotional information and semantic information. This separation has been argued on psychological and neural grounds (e.g. Adolphs et al. 2002). It has led to a state of affairs in which emotions are investigated in their non-linguistic manifestation (e.g. expression and recognition via face and voice), whereas language referring to emotions is considered to be less interesting, being considered secondary to the basic non-linguistic expressions of emotions. Recent work, however, has begun to challenge this tradition, suggesting there is mutual interaction between the affective and cognitive systems (see Dolan 2002 for a review). For example, negative words presented subliminally are better identified than neutral words (Gaillard et al. 2006) and trigger long lasting effects in the amygdala (Naccache et al. 2005). These results indicate that the primarily subcortical system engaged in processing emotion from non-verbal stimuli (i.e., faces) is also engaged in processing emotional valence in words. This suggests interactions between language processing and the limbic system along similar lines as it has been argued above for sensory-motor systems, thus, supporting the idea of a foundational role for affect. In tasks such as lexical decision and word naming, processing differences between emotionally valenced and neutral words are abundant in the literature. In a number of studies, a slowdown for negatively valenced words has been reported (e.g. Algom et al. 2004, Estes and Adelman 2008), although other studies found an advantage in the processing of emotionally valenced words, regardless of polarity (e.g. Eviatar and Zaidel 1991; Kanske and Kotz 2007).

In our own work, we have found that valence has a facilitatory role in word recognition (Kousta et al. 2009). In a lexical decision experiment in which we manipulated valence and controlled for a large number of lexical and sublexical variables, as well as arousal, we found that emotionally valenced words, regardless of polarity, were recognised faster than neutral words (we collected valence and arousal norms for 1,200 words following

the procedure used in the ANEW database, Bradley and Lang 1999). Moreover, we replicated this finding in a regression analysis of lexical decision RTs for 1,446 English words (taken from the English Lexicon Project, ELP, Balota et al. 2007), in which we used as predictors the same variables considered in the experiment. We found that previous conflicting results reported in the literature can be accounted for either in terms of limited control of lexical and sublexical variables (as also shown by Larsen et al. 2006) or in terms of characteristics of the distribution of valence ratings (unimodal vs. bimodal) used in regression analyses (this is the case for the results reported by Estes and Adelman 2008). This result is important for our purposes because it demonstrates that emotional valence (regardless of polarity) affects word recognition and that this is the case not just for emotion words but for all emotionally-valenced words. In general terms, this result is compatible with a motivational model of affective states as proposed by Lang and colleagues (Bradley and Lang 2000; Lang et al. 1990) proposing that emotion-evoking stimuli are motivationally relevant and that they lead to rapid modification of behaviour.

Whereas this work sets the stage for a role of affective valence (regardless of polarity) in semantic processing, importantly, in related work, we have provided initial evidence that valence may play a greater role in learning and representing abstract than concrete words (Kousta et al. submitted). First, as mentioned above, it is clearly the case that emotion words are abstract; further, many emotionally valenced words are also abstract. Why would this association matter? We argue that affect may help in the learning of abstract words that would otherwise be more difficult to learn because of their lack of sensory-motor associations. Some suggestion in favor of such a possibility comes from a regression analysis that we carried out in which we found that affect is a significant predictor of the age in which abstract words are acquired. We partitioned the concreteness scale at the mean and using Age of Acquisition ratings from the Bristol norms (Stadthagen-Gonzalez and Davis 2006) we regressed age of acquisition ratings on valence ratings and found that they are related by a U-shaped function. Although valence explains just under 6% of the variance in AoA ratings, we find appealing the possibility that emotion may provide a bootstrapping mechanism for the acquisition of abstract words. The relation between AoA and valence rating is reported in Figure 2a.

More critically, Kousta et al. (submitted) reported a further finding that supports the hypothesis of a differential role of affective information in the representation of concrete and abstract words. In lexical decision experiments, as well as in regression analyses of data from the ELP, we have found that after controlling for *imageability* (which is a primary

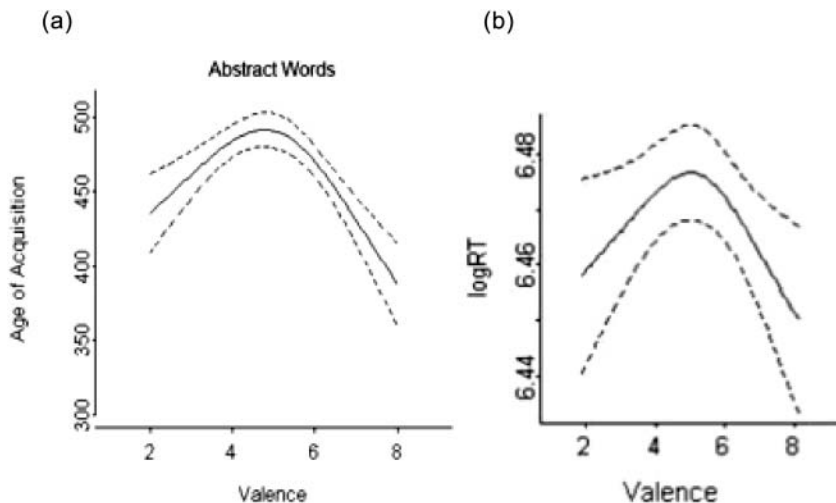


Figure 2. Plots of partial effects of valence (1 = negative; 5 = neutral; 9 = positive). Dashed lines indicate 95% confidence intervals. (a) Effect of valence on age of acquisition for abstract words (the concreteness scale was partitioned at the mean. Age of Acquisition ratings from Stadthagen-Gonzales et al. 2006), taken from Kousta et al. submitted. (b) Partial effect of valence as predictor of lexical decision reaction times for 1,446 words, once a large number of lexical and sublexical factors, see footnote 1 are partialled out (RT data from Balota et al. 2006), reprinted from *Cognition* 112(3), Stavroula-Thaleia Kousta, David P. Vinson and Gabriella Vigliocco, *Emotion words, regardless of polarity, have a processing advantage over neutral words*, pages 473–481, Copyright (2009), with permission from Elsevier

determinant of the difference between concrete and abstract words according to the dual coding theory by Paivio 1986), and *context availability* (namely the number of contexts in which a given word can be used, considered to be the primary difference between concrete and abstract words according to the context-availability hypothesis by Schwanenflugel 1991) in addition to a large number of other lexical and sublexical factors we found that abstract words have a processing advantage over concrete words. This surprising result contrasts with the previous literature in which a processing advantage for concrete over abstract words has been consistently replicated. Importantly in previous work, imageability (and often familiarity) was not controlled; thus the previous results show that words that are more imageable, and therefore have greater sensory-motor associations, are processed faster than words that are less imageable, but not that concrete words have an advantage over abstract words.

We found that this residual advantage for abstract over concrete words may be accounted for in terms of differences in the *valence* of the words. Most relevant here are the findings from regression analyses of RT data from the ELP for a set of 1,446 words for which concreteness, imageability, familiarity, AoA and valence norms are available. We carried out two sets of analyses: one in which all the variables above (including all other lexical and sublexical factors) were included but not valence, and one in which valence was added as a predictor. When valence was not included in the model, both concreteness and imageability were significant predictors of lexical decision reaction times. In the model with valence added, we found that valence was a significant predictor of latencies (Figure 2b shows the partial effect of this variable), predicting inhibition for neutral words and facilitation for emotional words. However, neither concreteness nor imageability were significant predictors of latencies in this model. This result strongly suggests an important role for affect in processing abstract words.

Below we discuss the role of linguistic information and how language and experience can be integrated in semantic representation.

4. Evidence for a role of linguistic information

It is intuitively obvious that we learn the meaning of a large number of words either implicitly or explicitly through language. This happens, for example, when we encounter words in texts (e.g. dictionaries or textbooks) or have the meaning of a word explained to us verbally. Thus, language must play an important role in learning word meanings. It is indeed the case that traditional *amodal* theories of semantic representation postulate some sort of language internal representation of word-meaning. This may be easily illustrated by considering semantic network models (e.g. Collins and Loftus 1975) according to which, semantic effects in behavioural tasks come about as a consequence of type of and strength of associative links between lexical concepts (see Vigliocco and Vinson 2007 for a discussion).

The linguistic nature of semantic representations for abstract words is argued in models such as the dual coding theory of Paivio (1986) according to which, abstract words would be represented solely in terms of a “verbal code”¹. This claim finds partial support in the imaging literature, given that the majority of studies that have compared processing of

1. Note that Paivio’s proposal does not make clear what a verbal code is. In particular whether the emphasis is on linguistic *form* or on linguistic *meaning*.

concrete/more imageable to abstract/less imageable words using tasks that require automatic access to semantics have reported greater activations for abstract words in the left language network (including left inferior frontal gyrus, IFG and left superior temporal sulcus, STS) (see Sabsevitz et al. 2005 for reviews) a result that has been argued to be compatible with dual coding and therefore with proposals (such as ours) that emphasise a role for linguistic information in semantic representation of especially abstract words. It is the case, however, that by no means there is agreement in the literature with respect to a greater engagement of the language network in the processing of abstract rather than concrete words. Critically, it is important to note that all imaging studies have confounded concreteness and imageability. However, as we have already discussed, the two constructs can be, and should be distinguished. Roughly speaking, concreteness is an operationalization of the distinction between entities and events that exist in the physical world and entities and events that exist in the human mind; imageability is an operationalization of the relevance of especially visual sensory properties of entities and events.

The importance of linguistic information in building semantic representations is also emphasized by theories that propose that word meanings are learned from how words are used in the language. Firth (1957), for example, suggested that “You shall know a word by the company it keeps” and that we learn at least part of the meaning of a word from “its habitual collocation” with other words. Along similar lines, Harris (1954) proposed that “if (two words) A and B have almost identical environments . . . we say they are synonyms” while “if A and B have some environment in common and some not . . . we say they have different meanings with the amount of meaning difference corresponding roughly to the amount of difference in the environments”. These ideas were to inspire the work of Lund and Burgess (1996) who developed the *Hyper-space Analogue of Language* (HAL) and of Landauer and colleagues who developed the *Latent Semantic Analysis* (LSA) and more recently, the work by Griffiths and colleagues (e.g. Griffiths et al. 2007) who developed a probabilistic *topic* model of semantic representation. In this latter model (that can be considered to be a probabilistic generalization of LSA) each text in a corpus can be seen as a probabilistic weighting of a set of discourse topics, with each discourse topic corresponding to probability distributions over words that emphasize a given theme. For example, a discourse topic labelled *sport* may place most of its probability mass on words like *game, ball, play, team, competition etc.* In these models, learning is the process of inferring the component topics. Each word in the model can then be represented as a distribution over these latent-

topics and this can be taken to be its semantic representation. Example of topics (taken from Andrews et al. in press) are provided in Table 3b.

The plausibility of inferring at least some aspects of word meanings from language, as exemplified in these models, has been established in a number of studies that have shown how measures of semantic similarity derived from these models can predict to some extent semantic phenomena in language, with the more recent developments (the topic model) outperforming previous models such as HAL and LSA. For example, Griffiths et al. (2007) show that the topic model does perform well in predicting association norms, priming data and reading times.

One virtue of models based on linguistic information is their ability to capture semantic representation across all domains with only one principle: they take into account only one type of information as contributing to semantic representation, i.e. distributional linguistic information, and they use only a single operation to learn semantics across domains, i.e. co-occurrences in text. Thus, they are parsimonious. They have, however, notable limitations. One of the most important is that these models are ungrounded and disembodied. In other words, linguistic information can only describe the relationship of words to one another but not to the physical world, or anything else beyond language itself. This is known as the symbol grounding problem (Harnad 1990). For example, Glenberg and Robertson (2000) reject models of semantic representation that are based solely on linguistic information, arguing that “to know the meaning of an abstract symbol such and English word, the symbol has to be grounded in something other than more abstract symbols”. Even if one were not to fully agree with criticisms to models of this type as those by Glenberg and Robertson (2000) or Harnad (1990), it is the case that because of the disembodied nature of representations in models such as HAL, LSA or the Topics model, they cannot account for any of the behavioural and neuroscientific evidence we have discussed above showing embodiment effects. In these models, although two words may have similar distributional patterns, and from that we can infer that they are semantically related, what they refer to in the world is unknown.

This fundamental criticism has led some researchers to consider these distributional patterns in the language simply as *epiphenomenal* and hence to argue that linguistic information does not play any role in meaning representation (e.g. Glenberg and Robertson 2000). As we discuss below, however, it is undeniable that this information is available to the learner, and that as such an optimal learning strategy is to avail of it. Of course, the fact that the information is available in great quantity does not *ipso facto* imply that it is actually used. However, we provide computational evidence that support the hypothesis that linguistic information is, in

fact, used in semantic representation in combination with experiential information.

5. Combining experiential and linguistic information

The idea that semantic representation comes about as a combination of both experiential and linguistic information is not unique to our proposal. Dual coding theory by Paivio (e.g. Paivio 1986) and also the multiple semantic stores hypothesis by Warrington and colleagues (e.g. Warrington and McCarthy 1994) make related claims that semantics is given by some combination of sensory-motor (primarily visual in these theories) and linguistic information. Moreover, dual coding theory and the multiple semantic stores ideas further assume that concrete and abstract words differ; whereas concrete words would be represented both in the visual and verbal store, abstract words would be represented only in the verbal store. More recently, Barsalou and colleagues (e.g. Barsalou et al. in press) have also spelled out a theoretical proposal (the Language and Situated Simulation framework, LASS) in which both linguistic and situated information crucially contribute to semantics (more precisely, in their view they both contribute to conceptual representation). This latter proposal converges and is complementary to the view we propose here. It is convergent in that both in our view and in LASS experiential/situated information are qualitatively different from linguistic information. It is complementary because whereas our work thus far has focused on demonstrating how these two types of information can be combined using statistical principles, work carried out by Barsalou and colleagues has focused on mapping the different time course of engagement of linguistic and situated information (e.g. Simmons et al. 2008).

Let us now describe our computational work. Andrews et al. (2009) used Bayesian modelling in order to develop and compare models of semantic representation that are based on the same architecture but learn from different types of data: either experience-only, or language-only, or a combination of the two. Crucial to this work is the data used as proxy of experience and language. With respect to experience, following a long standing tradition in cognitive psychology and neuropsychology (e.g. Farah and McClelland 1991; Hinton and Shallice 1993; McRae et al. 1997; Rosch and Mervis 1975; Vigliocco et al. 2004), we used speaker-generated features as a proxy to the types of experiential properties most relevant for different concepts. Speaker-generated features are typically obtained by asking subjects to produce as many properties as possible that together define and describe a given concept. Speaker generated features are not the only way to tap into the experiential information most

Table 1. *Sample of speaker-generated features for two nouns referring to objects and two verbs referring to actions. Weights (in brackets) reflect the number of speakers (max = 20) who generated that feature for a given word (from Vinson and Vigliocco 2002)*

the-grapefruit	the-hatchet	to-blink	to-pound
fruit (18)	sharp (14)	use-eye (16)	hit (14)
pink (17)	cut (13)	close (14)	beat (8)
yellow (13)	tool (13)	involuntary (11)	hard (6)
sour (11)	wood (9)	open (8)	force (5)
juice (8)	axe (8)	action (7)	noise (5)
eat (7)	chop (7)	protect (7)	use-fist (5)
breakfast (6)	handle (7)	fast (5)	anger (4)
round (6)	metal (7)	move (5)	action (3)
citrus (4)	blade (6)	use-eyelid (5)	contact (3)
large (4)	small (4)	reflex (4)	flat (3)
orange (4)	for-humans (3)	intentional (3)	loud (3)
food (3)	weapon (3)	sudden (3)	move (3)
seed (3)	danger (2)	by-animal (2)	object (3)
sweet (3)	survive (2)	by-humans (2)	physical (3)
bitter (2)		natural (2)	punch (3)
healthy (2)			use-hammer (3)
			violent (3)

relevant for different words and they are susceptible to criticisms related to the linguistic nature of the feature collection task, but we take them to be, nonetheless, the best proxy available to us. Their usefulness in investigating experiential (modality-related) effects in semantic representation has been demonstrated in a number of studies (Cree and McRae 2003; McRae et al. 1997; Vigliocco et al. 2002; Vigliocco et al. 2004; Vinson et al. 2003). Table 1 provides examples of the features produced.

With respect to linguistic information, following a long standing tradition in cognitive science and computational linguistics (Schutze 1992; Burgess and Lund 1997; Landauer and Dumais 1997; Griffiths et al. 2007), we used a corpus of text, taken from British National Corpus (BNC).

We developed three main models: the experiential model, the language-based, or distributional, model and the combined model. In the experiential model, words are defined as probability distributions over latent feature classes where these latent feature classes are in turn distributions over the speaker-generated features. In the distributional model, texts are latent distributions over discourse-topics, where these topics are probability distributions over words. The combined model is realised exactly on the same principles; however, here texts are distributions over latent

Table 2. *Examples of randomly chosen latent distributions of features in the experience model (a); words in the language model (b) and both features and words in the combined model. In each example, each column represents a topic. For each topic, we report the most highly probable features, words and both features and words (see Andrews et al. 2009)*

A)						
juice	fur	speak	wheel	mix	construct	leg
yellow	4-legs	word	transport	rotate	build	fast
red	tail	voice	passenger	spoon	new	exercise
round	pet	talk	gas	turn	fix	feet
grow	big	mouth	automobile	utensil	work	slow
sweet	small	language	drive	dance	create	body
sour	bark	sound	metal	hand	building	intentional
green	black	express	seat	bowl	material	go
taste	hair	converse	door	repeat	house	upright
seed	farm	noise	window	join	break	walk
small	wild	secret	sport	combine	form	sport
peel	domestic	explain	fast	stretch	physical	arm
citrus	ear	say	engine	awkward	action	face
good	white	comfort	move	cook	hole	foot
skin	ride	verbal	destination	bake	hand	speed
eat	zoo	understand	large	stir	finish	shoe
orange	wool	friend	oil	kitchen	carpenter	sweat
pit	friend	gossip	expensive	container	wood	destination
soft	meow	command	plastic	liquid	heavy	long
flesh	large	share	low	repeat	machine	work
B)						
league	prison	rate	pub	market	railway	air
cup	years	cent	guinness	stock	train	aircraft
season	sentence	inflation	beer	exchange	station	flying
team	jail	recession	drink	demand	steam	flight
game	home	recovery	bar	share	rail	plane
match	prisoner	economy	drinking	group	locomotive	airport
division	servicing	cut	alcohol	news	class	pilot
win	office	fall	bottle	trading	run	fly
club	life	economic	whisky	following	engine	jet
games	appeal	year	spirits	index	track	crash
final	case	rise	brewery	yesterday	lines	near
play	justice	confidence	wine	close	running	aviation
home	escape	industry	pint	fall	valley	base
won	cell	billion	brewing	early	passenger	airline
football	given	yesterday	drunk	added	service	flew
coach	guilty	growth	real	stocks	built	ground
second	punishment	spending	ale	value	platform	crew
victory	judge	high	duty	better	freight	squadron
saturday	term	increase	lager	fell	great	helicopter
round	crime	sales	cider	london	british	force

Table 2. (Continued)

C)							
W	mouth	instruct	transport	body	food	need	sound
	liquid	knowledge	wheel	hand	cook	give	ear
	consume	learn	vehicle	joint	combine	money	listen
	food	give	gas	move	pot	purchase	sense
	swallow	school	car	arm	heat	own	noise
	ingest	talk	passenger	humans	hot	trade	hear
	enjoy	information	small	connect	eat	return	soft
	hunger	show	automobile	muscle	stir	borrow	vibrate
	taste	help	fast	bone	oven	goods	comprehend
	thirst	teach	adult	bend	cut	swap	music
	stomach	express	move	point	stove	bank	word
	action	experience	destination	finger nail	utensil	buy	voice
	nutrition	task	metal	part	mix	store	brain
	water	effort	door	limb	bake	service	yell
	chew	authority	drive	up	pan	take	communicate
	body	guide	humans	reach	liquid	lend	perceive
	survive	idea	steer	finger	kitchen	gift	involuntary
	digest	grow	seat	touch	knife	business	physical
	fat	allow	engine	leg	prepare	generous	process
	daily	lead	window	shoulder	water	interest	pitch
	F						
	food	course	car	arms	add	bank	heard
	eat	students	road	arm	cook	exchange	hear
	drink	english	drive	fingers	oil	loan	hearing
	eating	language	driving	side	minutes	loans	sound
	wine	education	cars	hands	chopped	lend	listen
	drinking	college	driver	shoulder	heat	mortgage	sounds
	drinks	university	drove	body	serve	borrow	listening
	alcohol	teaching	van	shoulders	large	terms	ear
	ate	student	vehicle	knee	butter	exchanged	noise
	meal	taught	front	wrist	salt	banks	listened
	lunch	courses	vehicles	elbow	mix	interest	ears
	weight	teach	engine	leg	pan	deal	loud
	diet	study	speed	finger	stir	borrowed	speak
	sugar	higher	drivers	chest	tbsp	purchase	mouth
	wine	learn	motor	slowly	sauce	finance	word
	foods	learning	lorry	exercise	mixture	cheque	call
	family	knowledge	traffic	knees	cream	accounts	tone
	meals	library	wheel	leg	oz	figure	full
	finished	training	passenger	pain	sugar	world	voice
	swallowed	cambridge	park	lower	egg	money	hears

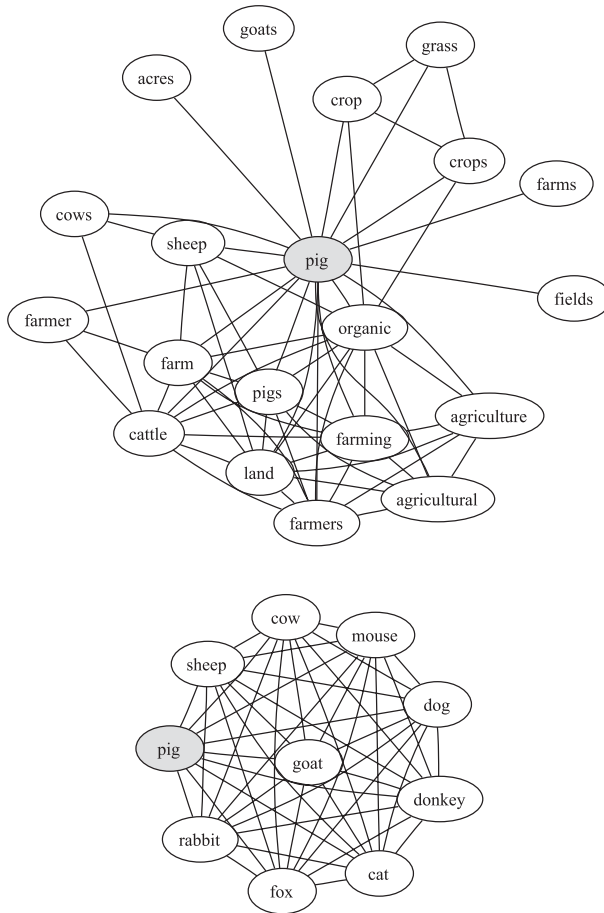


Figure 3. *Neighborhood cliques of kill according to the experience, language and combined models*

semantic classes where these classes are probability distributions over *both* features and words. The objective of the models is to discover latent statistical patterns in the training data-sets. For the experience model, these are clusters of correlated features, for the language model, these are clusters of correlated words, and for the combined model the clusters capture correlations both within and between data-sets. Table 2 presents examples of the latent distributions developed by each model.

Figure 3 presents the neighbourhood of the word *kill* in the experience, language and combined models. This example illustrates how the experi-




Figure 3. (continued)

ence and language models pick up on somewhat different senses of *kill*, hence providing qualitatively different semantic information. These are then integrated (not simply summed up) in the combined model which develops latent semantic classes on the basis of the correlation within features, text and between them.

Quantitative comparisons of correlations comparing the neighbourhood structures in each model further support the qualitatively different nature of the representations developed by the experience and language model showing that the correlations between the experience and language model is low (see Andrews et al. in press for further details).

Next, we carried out model comparisons to assess how good the different models are in accounting for semantic effects in behaviour. These




Figure 4. *The HPD regions of the mean rank percentile score for each model with respect to the Nelson's and the EAT association norms (adapted from Andrews et al. 2009)*

comparisons were carried out on three main types of data: association norms (both from Nelson et al. (1996) as well as the Edinburgh Association norms); semantically related word substitution errors (unpublished corpus made available to us by T. Harley), and data from semantic priming experiments (McRae et al. 1997 and Vigliocco et al. 2004). These different data sets cover semantic similarity measures taken from off-line judgements (association norms); but also on-line production (substitution errors) and comprehension (priming) semantic effects.

In addition to the three models just discussed, we included another model relevant here. In this additional model (the “independent model”), semantic similarity measures from the experience and from the language models were averaged to assess whether taking into account correlations between experience and language data (as in the combined model) truly provides something more than the sum of the parts. For illustrative purposes, in Figure 4, we present the results of the model comparison for the association norms. Comparisons using the other datasets showed the same pattern. The figure presents the high posterior density (HPD) region of the mean rank percentile score for each model with respect to the Nelson and in the EAT norms. As it can be seen, all models do quite well, but most importantly the combined model outperforms the others, ranking first.

Thus to summarise, our computational work provides evidence that (a) experience and language do provide qualitatively different information

Table 3. *Inferred or predicted features for concrete (a) and abstract (b) words in the combined model.*

A) Concrete words

accident	army	bowl	cigar	harbour	alcohol	rat
blood	attack	soup	smoke	sail	drink	fur
flesh	war	spoon	scent	boat	thirst	wild
life	destroy	stove	odor	ship	ingest	hop
kill	anger	pan	disgust	water	consume	cute
death	kill	oven	emit	float	liquid	long
drive	military	heat	breath	lake	swallow	whisker
wheel	oppose	carrot	air	ocean	enjoy	tail
passenger	explode	bake	inhale	cargo	glass	pet
Vehicle	deadly	silverware	exhale	transport	mouth	squeak
Window	pain	cook	gross	steer	stomach	teeth

B) Abstract words

obsession	feeling	sensation	chaos	luxury	fashion	death
women	heart	heart	demolish	comfort	women	sad
scream	touch	touch	attack	beautiful	grace	sudden
crave	rough	rough	explode	texture	beautiful	sick
desire	love	physical	war	cloth	style	harsh
relieve	sense	thirst	collide	gift	wear	black
love	texture	sense	impact	decorate	casual	cold
discomfort	hot	stomach	pain	gold	leather	dark
burn	finger	ingest	destroy	desire	cloth	bad
explode	cold	healthy	deadly	enjoy	cotton	violent
attack	hand	love	force	art	party	blood

(and, hence, renders the idea that the distributional patterns found in language are simply epiphenomenal less plausible) and (b) combining experience and language provides us a better approximation to semantic effects in behavioural tasks. Thus, this work provides critical evidence for the hypothesis that semantic representation comes about as a statistical combination of information from experience (in particular, sensory-motor and affective) and linguistic information. Moreover, by being combined with experiential meaning, aspects of meaning learnt from linguistic data are not disembodied but become hooked up to the world.

This latter fact is important. In the combined model, inferences concerning the likely experiential features of words learnt only via language can be made. This parallels the situation in which a child has experienced, for example, some animals in her/his physical environment and then (s)he learns about other animals from reading stories. Table 3a provides examples of this situation (in the models we developed these are cases where a

semantic representation was developed only on the basis of texts as no speaker-generated features were available for those words). As can be seen, these inferences are not always correct, for example “pet” is inferred as a feature of “rat”, but on the whole they are nonetheless statistically plausible. These inferences are obviously not limited to concrete words, Table 3b provides examples of inferred features for some abstract words. It is an empirical question for further research to assess whether inferences of this type, including errors such as those made by the model, are in fact typical of vocabulary learning during development.

6. Conclusions

We have presented here a general framework for how semantics, the meaning of words, is represented and learnt. Our approach is in line with general embodied views of cognition in stating a central role for sensory-motor systems: language processing automatically and immediately engages sensory-motor representations shared with perception and motor control (e.g. Barsalou et al. in press; Pulvermüller et al. 2005). Semantics is not only embodied in externally derived sensory-motor representations. We also argue that once we move from concrete to more abstract domains of knowledge, language processing automatically and immediately engages the system that processes emotions. Thus, semantics is grounded into our interactions with the external environment and also in our experience of our inner states. This grounding is fundamental because if the language system could not take advantage of already formed sensory-motor and affective representations, word meanings would not be learnable. After word learning has first been grounded in experiential information, linguistic information can and does provide another extremely rich source of data from which meaning can be learnt. Importantly the use of linguistic information not only enriches our semantic knowledge of already existing representations, but also allows us to speed up learning of new vocabulary for which experiential properties can be inferred in the absence of direct experience.

There are aspects of our view, which are currently more supported by the evidence (sensory-motor grounding); others which are still more speculative (differential role of affect for concrete and abstract words). It is also the case that a lot more computational work is needed to establish how more fine-grained sequential and syntactic linguistic information contributes to meaning representation. We believe, however, that the current proposal has the virtue of forcing us to consider semantic learning and representation from a somewhat different perspective, moving beyond the simplified divide between embodied and amodal theories.

References

- Adolphs, R., S. Baron-Cohen & D. Tranel. 2002. Impaired recognition of social emotions following amygdala damage. *Journal of Cognitive Neuroscience* 14. 1264–1274.
- Algom, D., E. Chajut & S. Lev. 2004. A rational look at the Emotional Stroop phenomenon: A generic slowdown, not a Stroop effect. *Journal of Experimental Psychology: General* 133. 323–338.
- Altarriba, J., & L. M. Bauer. 2004. The distinctiveness of emotion concepts: A comparison between emotion, abstract, and concrete words. *American Journal of Psychology* 117. 389–410.
- Andrews, M. A., D. P. Vinson & G. Vigliocco. 2009. Integrating experiential and distributional data to learn semantic representations. *Psychological Review* 116. 463–498.
- Balota, D., M. Cortese, K. Hutchison, B. Kessler, B. Loftis, J. Neely, D. Nelson, G. Simpson, R. Treiman & M. Yap. 2007. The English Lexicon Project. *Behaviour Research Methods* 39(3). 445–459.
- Barsalou, L. W. 1999. Perceptual symbol systems. *Brain and Behavioural Sciences* 22. 577–660.
- Barsalou, L. W., A. K. Barbey, W. Kyle Simmons & C. D. Wilson. 2003. Grounding conceptual knowledge in modality-specific systems. *Trends in Cognitive Sciences* 7(2). 84–91.
- Barsalou, L. W., A. Santos, W. K. Simmons & C. D. Wilson. in press. Language and simulation in conceptual processing. In M. De Vega, A. M. Glenberg & A. C. Graesser (eds.), *Symbols, embodiment, and meaning*. Oxford: Oxford University Press.
- Bloom, L. 1998. Language acquisition in its developmental context. In D. Kuhn & R. S. Siegler (eds.), *Handbook of child psychology* 2, 309–370. New York: Wiley.
- Boroditsky, L. & M. Ramscar. 2002. The roles of body and mind in abstract thought. *Psychological Science* 13. 185–188.
- Boulenger, V., V. Deprez, M. Jeannerod, T. A. Nazir, Y. Paulignan & A. C. Roy. 2006. Cross-talk between language processes and overt motor behaviour in the first 200ms of processing. *Journal of Cognitive Neuroscience* 18(10). 1607–1615.
- Bradley, M. M. & P. J. Lang. 1999. *Affective norms for English words (ANEW): Stimuli, instruction manual and affective ratings* (Technical report C-1). Gainesville, FL: The Center for Research in Psychophysiology, University of Florida.
- Bradley, M. M. & P. J. Lang. 2000. Affective reactions to acoustic stimuli. *Psychophysiology* 37. 204–215.
- Bretherton, I. & M. Beeghly. 1982. Talking about internal states: The acquisition of an explicit theory of mind. *Developmental Psychology* 18. 906–921.
- Britten, K. H., J. A. Movshon, W. T. Newsome & M. N. Shadlen. 1992. The analysis of visual motion: A comparison of neuronal and psychophysical performance. *Journal of Neuroscience* 12. 4745–4765.
- Buccino, G., F. Binkofski, G. R. Fink, L. Fadiga., L. Fogassi, L. Gallese, V. Seitz, R. J. Zilles, G. Rizzolatti & H. J. Freund. 2001. Action observation activates premotor and parietal areas in a somatotopic manner: An fMRI study. *European Journal of Neuroscience* 13. 400–404.
- Buccino, G., L. Riggio, G. Melli, F. Binkofski, V. Gallese & G. Rizzolatti. 2005. Listening to action-related sentences modulates the activity of the motor system: A combined TMS and behavioural study. *Cognitive Brain Research* 24. 355–363.
- Burgess, C. & K. Lund. 1997. Modeling parsing constraints with high-dimensional context-space. *Language and Cognitive Processes* 12. 177–210.
- Bush, G., P. Luu & M. I. Posner. 2000. Cognitive and emotional influences in anterior cingulate cortex. *Trends in Cognitive Science* 4. 215–222.

- Collins, A. M. & E. F. Loftus. 1975. A spreading-activation theory of semantic processing. *Psychological Review* 82(6). 407–428.
- Coulson, S. 2000. *Semantic leaps: Frame-shifting and conceptual blending in meaning construction*. Cambridge, UK: Cambridge University Press.
- Cree, G. S. & K. McRae. 2003. Analyzing the factors underlying the structure and computation of the meaning of chipmunk, cherry, chisel, cheese and cello (and many other such concrete nouns). *Journal of Experimental Psychology: General* 132. 163–201.
- Crutch, S. & E. Warrington. 2005. Abstract and concrete concepts have structurally different representational frameworks. *Brain* 128. 615–627.
- Decety, J. & J. Grezes. 2006. The power of simulation: Imagining one's own and other's behaviour. *Cognitive Brain Research*. 1079. 4–14.
- Dolan, R. J. 2002. Emotion, cognition, and behavior. *Science* 298(5596). 1191–1194.
- Estes, Z. & J. S. Adelman. 2008. Automatic vigilance for negative words is categorical and general. *Emotion* 8. 453–457.
- Eviatar, Z. & E. Zaidel. 1991. The effects of word length and emotionality on right hemisphere contribution to lexical decision. *Neuropsychologia* 29(5). 415–428.
- Farah, M. J. & J. L. McClelland. 1991. A computational model of semantic memory impairment: Modality-specificity and emergent category specificity. *Journal of Experimental Psychology: General* 120. 339–357.
- Fodor, J. 1976. *The language of thought*. Sussex: Harvester Press.
- Fodor, J. A., M. F. Garrett, C. H. Parkes & E. C. T. Walker. 1980. Against definitions. *Cognition* 8. 263–367.
- Firth, J. R. 1957. A synopsis of linguistic theory, 1930–1955. *Studies in Linguistic Analysis*. 1–32. Special Volume, Philological Society.
- Gaillard, R., L. Cohen, S. Dehaene, A. Del Cul, L. Naccache & F. Vinckier. 2006. Nonconscious semantic processing of emotional words modulates conscious access. *PNAS* 103. 7524–7529.
- Gallese V. & G. Lakoff. 2005. The brain's concepts: The role of the sensory-motor system in conceptual knowledge. *Cognitive Neuropsychology* 22. 455–479.
- Gibbs, R. 1994. *The poetics of mind: Figurative thought, language, and understanding*. New York: Cambridge University Press.
- Gibbs, R. 2006. *Embodiment and cognitive science*. New York: Cambridge University Press.
- Gleitman, L. R., K. Cassidy, R. Nappa, A. Papafragou & J. C. Trueswell. 2005. Hard words. *Language Learning and Development* 1. 23–64.
- Glenberg, A. M. & D. A. Robertson. 2000. Symbol grounding and meaning: A comparison of high-dimensional and embodied theories of meaning. *Journal of Memory & Language* 43. 379–401.
- Glenberg, A. M. & M. P. Kaschak. 2003. The body's contribution to language. In B. H. Ross (ed.), *The psychology of learning and motivation*, vol. 43, 93–126. San Diego, CA: Academic Press.
- Glenberg, A. & V. Gallese. submitted. *Action-based language: A theory of language acquisition, comprehension and production*.
- Glenberg, A. M., L. Cattaneo & M. Sato. 2008. Use-induced motor plasticity affects the processing of abstract and concrete language. *Current Biology* 18. R290–R291.
- Graesser, A. C., P. L. Hopkinson & C. Schmid. 1987. Differences in interconcept organization between nouns and verbs. *Journal of Memory and Language* 26. 242–253.
- Green, B. F. 1961. Figure coherence in the kinetic depth effect. *Journal of Experimental Psychology* 62. 272–282.
- Griffiths, T. L., M. Steyvers & J. B. Tenenbaum. 2007. Topics in semantic representation. *Psychological Review* 114. 211–244.

- Harnad, S. 1990. The symbol grounding problem. *Physica* 42. 335.
- Harris, Z. 1954. Distributional structure. *Word* 10. 146–162.
- Hauk, O., I. Johnsrude & F. Pulvermuller. 2004. Somatotopic representation of actions words in human motor and premotor cortex. *Neuron* 41. 301–307.
- Hinton, G. E. & T. Shallice. 1993. Lesioning an attractor network: Investigations of acquired dyslexia. *Psychological Review* 98. 74–95.
- Huttenlocher, J. & F. Lui. 1979. The semantic organization of some simple nouns and verbs. *Journal of Verbal Learning and Verbal Behavior* 18. 141–179.
- Jackendoff, R. 2002. *Foundations of language*. Oxford: Oxford University Press.
- Kanske, P. & S. A. Kotz. 2007. Concreteness in emotional words: ERP evidence from a hemifield study. *Brain Research* 1138. 138–148.
- Kaschak, M. P., M. Aveyard, A. A. Blanchard, C. J. Madden, D. J. Therriault, R. H. Yaxley & R. Zwaan. 2005. Perception of motion affects language processing. *Cognition* 94. B79–B89.
- Kaschak, M. P., M. Aveyard, R. H. Yaxley & R. Zwaan. 2006. Perception of auditory motion affects language processing. *Cognitive Science* 30. 733–744.
- Kemmerer, D., J. G. Castillo, T. Talavage, S. Patterson & C. Wiley. 2007. Neuroanatomical distribution of five semantic components of verbs: evidence from fMRI. *Brain and Language* 107. 16–43.
- Kousta, S.-T., D. P. Vinson & G. Vigliocco. 2009. Emotion words, regardless of polarity, have a processing advantage over neutral words. *Cognition* 112. 473–481.
- Kousta, S., M. Andrews, G. Vigliocco & D. P. Vinson (submitted). *The representation of abstract word meanings. Why emotion matters.*
- Lakoff, G. & M. Johnson. 1980. *Metaphors we live by*. Chicago: University of Chicago Press.
- Lakoff, G. & M. Johnson. 1999. *Philosophy in the flesh: The embodied mind and its challenge to Western thought*. Basic Books.
- Landauer, T. & S. Dumais. 1997. A solutions to Plato's problem: The Latent Semantic Analysis theory of acquisition, induction and representation of knowledge. *Psychological Review* 104. 211–240.
- Lang, P. J., M. M. Bradley & B. N. Cuthbert. 1990. Emotion, attention, and the startle reflex. *Psychological Review* 97. 377–395.
- Larsen, R. J., D. A. Balota & K. A. Mercer. 2006. Lexical characteristics of words used in emotional Stroop experiments. *Emotion* 6. 62–72.
- Levinson, E. & R. Sekuler. 1976. Adaptation alters perceived direction of motion. *Vision Research* 16. 779–781.
- Levelt, W. J. M. 1989. *Speaking: From intention to articulation*. Cambridge, MA: MIT Press.
- Levelt, W. J. M., A. S. Meyer & A. Roelofs. 1999. A theory of lexical access in speech production. *Behavioral and Brain Sciences* 22. 1–38.
- Lund, K. & C. Burgess. 1996. Producing high-dimensional semantic spaces from lexical co-occurrence. *Behavior Research Methods, Instrumentation, and Computers* 28. 203–208.
- McRae, K., V. R. de Sa & M. S. Seidenberg. 1997. On the nature and scope of featural representations of word meaning. *Journal of Experimental Psychology: General* 126. 99–130.
- Meteyard, L., B. Baharami & G. Vigliocco. 2007. The perceptual grounding of motion verbs. *Psychological Science* 11. 1007–1013.
- Meteyard, L. & G. Vigliocco. 2008. The role of sensory and motor information in semantic representation: A review. In P. Calvo and A. Gomila (eds.), *Handbook of cognitive science: An embodied approach*. Oxford: Elsevier.

- Meteyard, L., B. Bahrami, G. Vigliocco & N. Zokaei. 2008. Now you see it: Visual motion interferes with lexical decision on motion words. *Current Biology* 18. R732–733.
- Naccache, L., C. Adam, M. Baulac, S. Clemencau, L. Cohen, S. Dehaene, R. Gaillard & D. Hasboun. 2005. A direct intracranial record of emotions evoked by subliminal words. *Proc. Natl. Acad. Sci. USA* 102. 7713–7717.
- Nakamura, K., R. Hanajima, N. Hara, S. Kouider, K. Sakai, Y. Takayama & Y. Ugawa. 2003. Task-guided selection of the dual neural pathways for reading. *Neuron* 52. 557–564.
- Nelson, D., C. McEvoy & T. Schreiber. 1996. *Word association norms*. <http://u3.usf.edu/FreeAssociation/>.
- Oliveri, M., A. Caramazza, C. Finocchiaro, M. Gangitano, A. Pascual-Leone & K. Shapiro. 2004. All talk and no action: A transcranial magnetic stimulation study of motor cortex activation during action word production. *Journal of Cognitive Neuroscience* 16(3). 374–381.
- Paivio, A. 1986. *Mental representations: A dual-coding approach*. Oxford: Oxford University Press.
- Pulvermüller, F., O. Hauk, R. J. Ilmoniemi & V. V. Nikulin. 2005. Functional links between motor and language systems. *European Journal of Neuroscience* 21. 793–797.
- Ridgeway, D., S. A. Kuczaj II & E. Waters. 1985. Acquisition of emotion-descriptive language: Receptive and productive vocabulary norms for ages 18 months to 6 years. *Developmental Psychology* 21. 9013–908.
- Rizzolatti, G & L. Craighero. 2004. The mirror-neuron system. *Annual Review of Neuroscience* 27. 169–192.
- Rosch, E. & C. B. Mervis. 1975. Family resemblance: Studies in the internal structure of categories. *Cognitive Psychology* 7. 573–605.
- Sabsevitz D. S., D. A. Medler, M. Seidenberg & J. R. Binder. 2005. Modulation of the semantic system by word imageability. *NeuroImage* 27. 188–200.
- Scase, M. O., O. J. Braddick & J. E. Raymond. 1996. What is noise for the motion system? *Vision Research* 36. 2579–2586.
- Schutze, H. 1992. *Dimensions of meaning*. IEEE Computer Society Press.
- Schwanenflugel, P. 1991. Why are abstract concepts hard to understand? In P. Schwanenflugel (ed.), *The psychology of word meanings*, 223–250. Hillsdale, NJ: Erlbaum.
- Simmons, K., S. B. Hamann, C. N. Harenski, X. P. Hu & L. W. Barsalou. 2008. *fMRI evidence for word association and situation simulation in conceptual processing*.
- Stadthagen-Gonzalez, H. & C. J. Davis. 2006. The Bristol norms for age of acquisition, imageability and familiarity. *Behavior Research Methods* 38. 598–605.
- Tettamanti, M., G. Buccino, M. Danna, V. Gallese, M. C. Saccuman & P. Scifo et al. 2005. Listening to action-related sentences activates fronto-parietal motor circuits. *Journal of Cognitive Neuroscience* 17(2). 273–281.
- Tsushima, Y., Y. Sasaki & T. Watanabe. 2006. Greater disruption due to failure of inhibitory control on an ambiguous distractor. *Science* 15(314). 1786–1788.
- Turner, M. & G. Fauconnier. 2000. Metaphor, metonymy, and binding. In A. Barcelona (ed.), *Metonymy and metaphor at the crossroads*. Berlin: Mouton de Gruyter.
- Vigliocco, G., M. F. Damian, D. P. Vinson & W. Levelt. 2002. Semantic distance effects on object and action naming. *Cognition* 85. B61–B69.
- Vigliocco, G., D. P. Vinson, W. Lewis & M. F. Garrett. 2004. The meanings of object and action words. *Cognitive Psychology* 48. 422–488.
- Vigliocco, G., J. Warren, S. Siri, J. Arciuli, S. Scott & R. Wise. 2006. The role of semantics and grammatical class in the neural representation of words *Cerebral Cortex* 16. 1790–1796.

- Vigliocco, G. & D. P. Vinson. 2007. Semantic representation. In G. Gaskell (ed.), *Handbook of Psycholinguistics*. Oxford: Oxford University Press.
- Vinson, D. P., G. Vigliocco, S. F. Cappa & S. Siri. 2003. The breakdown of semantic knowledge: Insights from a statistical model of meaning representation. *Brain and Language* 86. 347–442.
- Warrington, E. K. & R. A. McCarthy. 1994. Multiple meaning systems in the brain: A case for visual semantics. *Neuropsychologia* 32. 1465–1473.
- Wellman, H. M., M. Banerjee, P. L. Harris & A. Sinclair. 1995. Early understandings of emotion: Evidence from natural language. *Cognition and Emotion* 9. 117–149.
- Zwaan, R. 2004. The immersed experiencer: Toward an embodied theory of language comprehension. In B. H. Ross (ed.), *The psychology of learning and motivation*, vol. 44. San Diego, CA: Academic Press.