Representing meaning:

A feature-based model of object and action words

David Patrick Vinson

UCL

Declaration

I, David Patrick Vinson, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Acknowledgements

First and foremost, thanks to Gabriella Vigliocco who has been an excellent supervisor and colleague. My development as a scholar has been a direct product of her constant ideas and inspiration, and she played a very important role in all of the work reported in this thesis. Thanks also go to James Hall for encouraging me into the study of cognition, and John Theios for taking a chance on a graduate student who had not especially applied himself as an undergraduate. The speaker-generated features could not have been collected without the help of many student research assistants at Wisconsin who joined me in this incredibly tedious exercise. The work presented here also benefited greatly from the involvement of collaborators including Stefano Cappa, Markus Damian, Merrill Garrett, Pim Levelt, Will Lewis and Simona Siri; and all the members of the Language (Production/ Processing) Lab for all the great discussion and mutual support. I also thank all my friends and family for helping me stay relatively sane during this long process, and especially my mother who never tired of asking about my progress. Special thanks to my secondary supervisor Judit Druks for helping me come up with a no-nonsense timetable for completing this thesis, insisting that I stick to it no matter what, and providing me so many detailed and useful comments and suggestions on various drafts. Thanks also to Jenny Sherrill for proofreading and editing. Finally, many, many thanks to Amanda for love and support, and making me happy.

Abstract

The representation of word meaning has received substantial attention in the psycholinguistic literature over the past decades, yet the vast majority of studies have been limited to words referring to concrete objects. The aim of the present work is to provide a theoretically and neurally plausible model of lexical-semantic representations, not only for words referring to concrete objects but also for words referring to actions and events using a common set of assumptions across domains. In order to do so, features of meaning are generated by naïve speakers, and used as a window into important aspects of representation. A first series of analyses test how the meanings of words of different types are reflected in features associated with different modalities of sensory-motor experience, and how featural properties may be related to patterns of impairment in language-disordered populations. The features of meaning are then used to generate a model of lexical-semantic similarity, in which these different types of words are represented within a single system, under the assumption that lexical-semantic representations serve to provide an interface between conceptual knowledge derived in part from sensory-motor experience, and other linguistic information such as syntax, phonology and orthography. Predictions generated from this model are tested in a series of behavioural experiments designed to test two main questions: whether similarity measures based on speaker-generated features can predict fine-grained semantic similarity effects, and whether the predictive quality of the model is comparable for words referring to objects and words referring to actions. The results of five behavioural experiments consistently reveal graded semantic effects as predicted by the feature-based model, of similar magnitude for objects and actions. The model's fine-grained predictive performance is also found to be superior to other word-based models of representation (Latent Semantic Analysis, and similarity measures derived from Wordnet).

Table of Contents

Tables

Figures

Representing meaning: A feature-based model of object and action words

Chapter 1: Introduction

The aim of the present work is to provide a psychologically and neurally plausible model of lexical-semantic representations for words referring to concrete objects and actions and events by collating features of meaning generated by naïve speakers. These features can be viewed as an indication of semantic representations, and when combined across multiple individuals, should provide some general insights into the meanings of words. Employing a feature-based semantic theory will allow the generation of a model of lexical-semantic similarity, in which different types of words are represented within a single system. This model will then be tested against behavioural data, and its predictive power will be assessed against extant models of semantic representation. The inclusion of the domain of actions and events into the model is important and innovative because nearly all previous research on word semantics has focused upon words referring to concrete objects and entities only.

The representation and organisation of word meaning

Meaning is a centrally important aspect of language which lies at the heart of communication. In language production, speaking is a continuous process of selecting the words that best correspond to the meaning of a message the speaker wishes to express (Levelt, 1989). Similarly, comprehending language is a continuous process of attempting to discern the meaning of a speaker's or writer's message. Word meaning provides the core information upon which all communication is built, and similarity in the meanings of words is invaluable in serving communication, for example, in providing definitions of technical terms (e.g. "The word deictic means 'pointing' or 'showing.'", Ballard, Hayhoe, Pook & Rao, 1995, p.725). Such uses of similarity is not restricted to formal written or instructional materials, but also occur in numerous situations in which an interlocutor is not familiar with a particular word being used, a situation in which it is quite normal to produce similar alternatives for explanatory purposes.

Such consequences of similarity in meaning may not necessarily reflect the underlying representations that are automatically consulted in normal conversational situations, but come into play only in situations in which word meaning are explicitly being discussed, and thus reflect a speaker's intuition about language rather than language itself. A more convincing demonstration of the impact of semantic similarity arises in cases in which semantic similarity does not facilitate communication, but instead has counterproductive consequences during online language processing. In production, this is particularly evident in slips of the tongue, where a substitution between one word and another (related in meaning) can result in a sensible sentence whose meaning is very different than the speaker intended:

> "US President Gerald Ford toasted Egyptian President Anwar Sadat and 'the great people of Israel--Egypt, excuse me.' " Dell (1995), p.183

In instances like these, semantic similarity between words can have undesired consequences. On the other hand, in online comprehension processes, semantic similarity can facilitate word recognition. For example, in the lexical decision paradigm, response time to a target word is faster if it is preceded by a semantically-related word than if it is preceded by an unrelated word (Meyer & Schvaneveldt, 1971). Such priming may serve the purpose of speeding up lexical processing, or at least improving the efficiency of accessing the correct meaning during comprehension.

Semantic similarity effects, both of a facilitatory and of an interfering nature, have been extensively studied in behavioural and neuroscientific research. Systematic investigation of spontaneously occurring speech errors reveals that lexical substitution errors are among the most frequent type of slips of the tongue (e.g. Fromkin, 1973; Garnham, Shillock, Brown, Mill & Cutler, 1981). The precise nature of similarity between the target and error word seems to vary, depending upon the semantic domain, but is predictable within a given domain (e.g., errors involving the names of body parts tend to be physically close to the intended body part, Garrett, 1992). Garrett's analyses of speech errors shows that, for nouns, most substitutions involve category coordinates (for example, *shoulders/elbows*, *eyes/ears*). In contrast, for verbs, errors of antonymy (*remember/forget*) are frequent, while coordinates (*drink/eat*; *looks/sounds*) are much less so.

More evidence of the consequences of semantic similarity comes from studies of semantic interference in naming tasks. This was pioneered in the work of Stroop (1935), who presented participants with words referring to the names of colours, printed in various colours of ink, and asked them to name the ink colour whilst ignoring the word itself (e.g., given the word "RED" printed in green ink, to say the word "GREEN"). The meanings of the written words had severe consequences on naming the ink colour: longer latencies and higher error rates. In the picture-word interference paradigm (a variant of the Stroop task) in which participants are asked to name a picture while ignoring a simultaneously appearing written word, semantically-related words interfere with picture naming (e.g. Glaser & Düngelhoff, 1984; Schriefers, Meyer & Levelt, 1990; Vigliocco, Vinson, Lewis & Garrett, 2004; Vigliocco, Vinson & Siri, 2005). In other contexts, semantic similarity has a facilitatory effect, such as semantic priming in comprehension as mentioned above (Meyer & Schvaneveldt, 1971; reviewed in Neely, 1991; Vigliocco, Vinson, Arciuli & Barber, 2008).

Semantic similarity also has consequences for patients whose semantic knowledge has been disrupted following brain injury. Especially relevant here are category-specific deficits, a phenomenon where patients are selectively impaired in some categories of knowledge and spared in others. The dissociation between the domains of living and non-living entities is best documented (e.g., Basso, Capitani, & Laiacona, 1988; Farah, Hammond, Metha, & Ratcliff, 1989; Hillis & Caramazza, 1995; Moss & Tyler, 2000; Sacchett & Humphreys, 1992; Sartori & Job, 1988; Sheridan & Humphreys, 1993; Vinson, Vigliocco, Cappa & Siri, 2003; Warrington & McCarthy, 1987; see Caramazza & Shelton, 1998), but numerous different patterns of finer-grained dissociations have also been reported, including selective impairment for body parts (McKenna & Warrington, 1978), animals (Caramazza & Shelton, 1998), fruits & vegetables (Hart, Berndt & Caramazza, 1985) and medical terms (Crosson, Moberg, Boone, Gonzales Rothi & Raymer, 1997; see Rogers and Plaut, 2002 for a review). Some patients with impairments for living things may also show a deficit for other (non-living) categories such as musical instruments, materials and liquids (Borgo & Shallice 2001; Siri, Kensinger, Cappa, Hood & Corkin, 2002; Warrington & Shallice, 1984). Patterns of impairment and sparing in category specific impairments offer further evidence for the psychological reality of semantic similarity because many such impairments seem to disproportionately affect semantically related clusters of items; they also provide important constraints on accounts of semantic representation, as will be discussed later.

Theories of semantic representation

How is semantic similarity represented in the mind and brain? Before attempting to answer this question, it is first necessary to outline the different theoretical frameworks within which the meanings of words corresponding to concepts may be represented. Theoretical accounts concerned with the representation of word meaning are strongly linked with conceptual categorisation and equivalence classification. In particular, they focus upon how different exemplars of a concept can be treated as equivalents, and can be assigned the same lexical label in a language.

Important issues in semantic theory are identifying the content of word meaning, the organisation of word meaning, and how the link between referent and word can be characterised.

Distinguishing semantics from concepts. Before describing the various theoretical perspectives, it is necessary to discuss the difference between conceptual-level representations and semantic representations. It is relatively uncontroversial that word meaning (i.e., lexical-semantics) is grounded in conceptual knowledge. More difficult is the question if the two are distinguishable from each other. The closeness of semantic and conceptual representations is clearly demonstrated by brain imaging research that shows, for example, that primary motor areas are activated when speakers see or hear words or sentences referring to actions (e.g. Hauk, Johnsrude & Pulvermuller, 2004; Martin & Chao, 2001; Tettamanti et al., 2005; Vigliocco, Warren, Siri, Arciuli, Scott & Wise, 2006). Moreover, it is often assumed that word meanings are indistinguishable from conceptual knowledge (e.g. Humphreys, Price & Riddoch, 1999).The working hypothesis of the present work, however is that conceptual and semantic knowledge are distinct levels of representation, each with its own distinct organisation. The present work also assumes that speaker-generated features provide a window into fundamental aspects of non-linguistic conceptual representations (such as the modality by which different types of information are learnt and experienced), and that the meanings of words are organised along different principles.

A number of arguments have been made in favour of the distinction between conceptual and semantic representations. For example, it has been pointed out that speakers of a language have many more concepts than words. For example, "the actions of two people maneuvering for one armrest in a movie theatre or airplane seat" is a familiar concept for which no verbal label exists (Murphy, 2002, p.389)[1]. This "more concepts than words" argument requires some sort of semantic level of representation (or at least a purely lexical level at which individual words are represented in some way, distinct from those concepts that are not lexicalised in a language), but does not require a separate level of organisation of this level. For example, concepts could be organised according to meaning whether they are lexicalised or not, and the semantic level of representation could simply be mapped from the relevant concepts on a one-to-one basis. Such a representational framework can be seen in the WEAVER^{++} model of lexical retrieval in production (Levelt, Roelofs & Meyer, 1999) in which the meanings of lexical concepts (and of lemmas) are not interconnected (i.e., there is no local organisation at this level).

Another argument that has been presented in favour of the conceptual-semantics divide is the cross-linguistic differences in mapping between conceptual and linguistic domains (see Vigliocco and Filipovic, 2004; Vigliocco & Vinson, 2007 for discussion). For example, English speakers have different words for *foot* and *leg*, while Japanese speakers have a single word (*ashi*) which refers to both. Similarly, English has numerous verbs corresponding to different manners of jumping: *leap*, *hop*, *spring*, *bounce*, *caper*, *vault*, *hurdle* and so on, while Italian does not (Slobin, 1996a). Differences of this nature can even be seen in how spatial situations are realised in two closely related languages such as English and Dutch. While English has two terms, *on* and *in*, Dutch has three: *aan*, *in* and *op* (Bowerman & Choi, 2003). Under the assumption of identity between conceptual and semantic knowledge, these language differences would require that the speakers of different languages also have different conceptual representations -

[1] Although nothing prevents speakers of a language from coining or adopting a new term for any concept. For example, the situation described above has been labelled "elbonics", originally by comedian Rich Hall, who has published a number of collections of such "missing lexical items" (e.g., Hall, R., 1984, *Sniglets (Snig'lit : Any Word That Doesn't Appear in the Dictionary, But Should).* Collier Books).

the view known as linguistic relativity (e.g. Davidoff et al., 1999; Levinson, 1996; Lucy, 1992; Roberson et al., 2000; Sapir, 1921; Sera et al., 2002; Slobin, 1996a, b; Whorf, 1956). But this conclusion only applies if conceptual and semantic representations are one and the same. If they are distinct levels of representation, these findings can be accommodated in the same manner as the "more concepts than words" argument. Conceptual organisation would be the same across languages; the only differences lie in which concepts are lexicalised.

However, there is evidence that seems to require not only a distinction between conceptual and semantic levels of representation, but also different principles of organisation at these distinct levels. These come from some language-specific effects related to meaning, which appear to be limited to semantic representations as they are only present in tasks that require verbalisation, but not in nonverbal tasks (Brysbaert et al., 1998; Kousta, Vinson & Vigliocco, in press; Malt, Sloman, Gennari, Shi & Wang., 1999; Vigliocco, Vinson, Paganelli & Dworzynski, 2005; see Slobin, 1996b).

Brysbaert et al. (1998) tested for language-specific effects of the manner in which speakers' languages require them to produce number words, either in forms like "four-and-twenty" (Dutch) or "twenty-four" (French). Participants were asked to report the solutions of simple mental calculations (e.g. "20 + 4 = ?" or "4 + 20 = ?"), either verbally or by typing the numbers on a keyboard. Cross-linguistic differences were observed such that participants were faster in providing the answers when the addends were presented in an order that matched the language (e.g. "20+4" when the answer is expressed "twenty-four", or "4+20" when the answer is expressed "four-and-twenty"). However, this pattern was only observed when responses were made verbally. The differences disappeared when participants were asked to type their responses in digits. This suggested that these differences were a product of verbal encoding rather than cross-linguistic differences at a conceptual level related to arithmetic operations.

Relevant results also come from investigations of the relationship between syntactic properties and semantic representations. Vigliocco, Vinson, Paganelli and Dworzynski (2005) investigated the effects of grammatical gender of Italian words on semantic relatedness and found that Italian words referring to animals sharing grammatical gender were judged to be semantically more similar, and were more likely to replace one another in slips of the tongue than words that were of different gender. Thus, grammatical gender, a syntactic property of words had semantic consequences. Crucially for the argument here, the effects of grammatical gender disappeared in similarity judgements upon pictures, a task that is most likely to tap conceptual knowledge. The study therefore was able to demonstrate the separability of conceptual and semantic levels of knowledge, and their separate respective organisation. In a follow-up study, Kousta, et al. (in press) induced slips of the tongue in bilingual speakers of Italian (L1) and English (L2) who performed the same task in both of their languages on different days. This study was designed to assess whether the above effects of grammatical gender would also be observed in a bilingual's second language – a pattern of results that would be predicted if such effects arise at a conceptual level (and/or if conceptual and semantic levels of representation are identical or mapped on a one-to-one basis). Instead, the errors made by bilingual speakers were comparable to those made by monolingual speakers, at least where grammatical gender is concerned.[2] Grammatical gender was reflected more in the bilinguals' errors in Italian than their errors in English for the same pictures, and this was true even when phonological similarity of the target and error words was taken into account (Kousta et al., in press). This provides strong support for an informational distinction between conceptual and semantic representations, showing not only that these effects of grammatical gender are limited

[2] The only qualitative differences in the errors came from certain errors that were related to cross-linguistic phonological effects, such as mistakenly producing *horse* for *bear* in English (the Italian word for *bear* is *orso*).

to semantic representations rather than arising at the conceptual level, but also that they do not extend to the semantic representations of a speaker's second language.

It is important to note that many theories concerning the representation of meaning are concerned with conceptual representations rather than semantic representations, or contain the implicit assumption that concepts and semantics are one and the same. Nonetheless, their stances with respect to content, organisation, and links between words and referents are still relevant to any discussion of semantic representation, particularly as the model of meaning that will be presented here addresses both conceptual and semantic levels of representation.

<u>Classical view of meaning.</u> Due to the importance of meaning in language, it is no surprise that questions related to the meanings of words have captured the interest of scholars since antiquity. Early theories of meaning, often termed the "classical view" (Smith & Medin, 1981, for a review), were based upon the assumption that the meaning of a concept (represented by a word) is its definition - a set of necessary features which would include all exemplars of the concept and exclude all others. Such a view of meaning, often couched in terms of formal logic, have been pervasive since classical times (e.g. Aristotle's <u>Categories</u>, 350 BCE/1941), and dominated theorising through much of the 1900s (e.g. Cassirer, 1953; Bourne, 1970; Katz & Fodor, 1963). The classical view, however, fell under severe criticism (see Smith & Medin, 1981; Mervis & Rosch, 1981; Murphy, 2002; Wittgenstein, 1953/2001). Wittgenstein (1953) demonstrated the apparent impossibility of producing adequate definitions to encompass all the meanings of a word, using the much cited example of the concept *game*. "*Game*" evades attempts to define it, as games need not be competitive, nor have scores, nor involve multiple participants, nor any other property of a subset of games that comes to mind. Wittgenstein points out that this difficulty extends to most concepts, not just *game*. Work by Rosch and colleagues provided further arguments

against the classical view. They argued that category boundaries are fuzzy rather than sharply delimited (Rosch, 1973; see also Hampton, 1979, for behavioural evidence), and they further developed Wittgenstein's notion of "family resemblance" in which prototypical members of a category are those which have the most attributes in common with other members of the category (Rosch & Mervis, 1975). Although a few researchers of semantics continue to pursue theoretical approaches similar to the classical view, most famously, Jackendoff (1990; 1992; 2002) who continues to explore the possibility of semantic representations in terms of primitives, most theories have diverged from the classical approach.

Relational theories of meaning. A broad class of alternative theories to the classical view focuses on characterising semantic representations by investigating meaning relations between words, rather than attempting to dissect the meanings of individual words themselves. A pioneer in this approach was Charles Osgood (see Osgood, 1962; Osgood, May & Miron, 1975; Osgood, Suci & Tannenbaum, 1957; Snider & Osgood, 1969), who developed the method of "semantic differential" that quantifies semantic relations between words by asking participants to rate individual words on a variety of attitude scales (e.g. good-bad, strong-weak, tense-relaxed). Crucially, rather than claiming that ratings for a given word on these individual scales revealed the componential features of the word's meaning, Osgood and colleagues used the semantic differential scale responses to generate measures of psychological distance between words (where greater proximity reflects higher similarity between words). Rather than treating each scale as independent, Osgood et al. applied factor analysis to reduce a high-dimensional representation space (one dimension for each of the attitude scales) into one of lower dimensionality. This allowed them not only to obtain a measure of the overall similarity between pairs of words, but also to evaluate the dimensions along which the words differed. Three distinct dimensions were repeatedly observed in these kinds of studies--dimensions which, Osgood argued, were universal and allowed evaluation of any semantic space in any situation: evaluative scales (e.g., good-bad); power scales (e.g., strong-weak); and activity scales (e.g., active-passive). These scale labels are descriptive of the first three dimensions obtained from factor analysis (i.e. those which explain the most variance in the data) across the scales that entered the factor analysis, and can be applied to nearly all semantic domains. Although Osgood's approach remains in use today in contexts such as advertising and marketing where evaluative judgments are important, the dimensions of evaluation, power and activity characterise semantic representations only in the broadest of terms, and relate only to a number of limited, relatively abstract domains. A bird, for example, is surely more than a combination of its ratings on evaluative, power and activity scales, which have nothing to do with specific physical form, activities performed, habitat, diet, or any other information that is important for the meaning of *bird*.

Similarly to Osgood's approach, according to semantic field theory (Trier, 1931; see Lehrer, 1974; Kittay, 1987), semantic representations arise from relationships among the meanings of different words. Semantic fields are considered to be a set of words that are closely related in meaning. The meaning of a word within a field is determined in terms of contrast to other words within the semantic field. In contrast to Osgood's approach, semantic field theory does not focus upon identifying broad dimensions that apply universally across all concepts, but instead attempts to identify the principles of contrast applicable within a field. For example, the semantic field of colour words is distinguished by hue and brightness (Berlin & Kay, 1969), kinship terms by age, sex, degree of relation (Bierwisch, 1969), cooking terms by factors like heat source, utensils involved and materials cooked (Lehrer, 1974), and body parts by function and proximity (Garrett, 1992).

Network models of semantic representation, while very different from semantic differential scaling or semantic field theory, also belong to the class of relational approaches to semantics. Network models go beyond earlier relational theories by specifying the details of the semantic relationships among words, rather than simply describing their distribution across the distinguishing dimensions of meaning. Early network models were semantic networks in which words are represented as nodes, and semantic relationships are expressed by labelled connections between nodes (e.g. Collins & Loftus, 1975). In this approach, a word's meaning is expressed by the links it has to other words, which other words is it connected to, and what types of connections are involved. Of paramount importance for network-based theories is the type, configuration and relative contribution of the links that exist between words. Numerous alternative frameworks have been developed (see Johnson-Laird, Herrmann & Chaffin, 1984 for a review) which differ along these dimensions. Importantly, these models have in common a focus upon (explicit) intensional relations, and a necessity to explicitly designate those relations.

Perhaps the most extensive model which implements distinct representational themes is Wordnet (Miller & Fellbaum, 1991), a network model of the representations of a large number of nouns, verbs and adjectives in English. In Wordnet, "nouns, adjectives and verbs each have their own semantic relations and their own organisation determined by the role they must play in the construction of linguistic messages" (p.197). These relations and organisation are constructed by hand based on the relations that are believed to be relevant within a given class of words. For nouns, the most important relations are synonymy, hierarchy and part-whole relations. For verbs, relations of troponymy (hierarchical relations related to specificity in manner), entailment, causation and antonymy are important.

All the relational theories described above depend upon deciding which relationships are most relevant in representing meaning, and then deciding upon a manner of implementation. An entirely different relational approach, however, seeks to discover representations of words in terms of their relationship to other words without making any assumptions about the organisational principles involved. This approach can be found in global co-occurrence models such as Latent Semantic Analysis (LSA, Landauer & Dumais, 1997) and Hyperspace Analogue to Language (HAL, Burgess & Lund, 1997). These models use large corpora of texts, computing aspects of a word's meaning based on other words found in the same linguistic contexts under the assumption that words that tend to share the same linguistic contexts will be similar in meaning. The resulting representations remain purely abstract, denoting a word's similarity to other words without revealing which aspects of meaning contribute to the observed similarity. Measures of similarity based on these models have been demonstrated to predict behavioural performance to some extent (see Burgess & Lund, 1997; Landauer & Dumais, 1997) suggesting that abstract relational representations derived from words' contexts (e.g. LSA and HAL) reflect patterns of similarity that have psychological plausibility.

Relational theories, however, have in common a serious flaw in that they focus only upon relationships among words and are not grounded in perception and action. As Johnson-Laird et al. (1984) wrote, "The meanings of words can only be properly connected to each other if they are properly connected to the world" (p. 313). Although Johnson-Laird referred to semantic network models, his criticism is relevant to any theory of representation that is not embodied in experience, at least, to some extent. It is largely with this concern in mind that many researchers developed perspectives that owe much to the classical view.

Featural theories of meaning. Although severe criticisms have been applied to the classical view, its general assumption that word meaning is componential in nature offers a way to ground meaning in perception and action. Different featural theories of meaning (e.g. Rosch & Mervis, 1975; Smith, Shoben & Rips, 1974; Collins & Quillian, 1969; Jackendoff, 1990; Minsky, 1975; Norman & Rumelhart, 1975; Shallice, 1993; Smith & Medin, 1981) consider the representation of meaning in terms of feature lists - those properties of meaning which, taken together, express the meaning of a word in some way. One class of these theories can be described as a modified version of the classical view with the incorporation of additional assumptions to avoid the criticisms aimed at it (Miller & Johnson-Laird, 1976; Smith & Medin, 1981). These accounts assume the existence of definitions for concepts in the classical sense (core features), but also another set of relevant features. These features (nonnecessary features) reflect information that is not necessarily part of the definition itself, but instead, properties of some, but not all, exemplars of a category. Although core features will always be relevant (because they are common to all members of a category), nonnecessary features would be used for identification procedures, as they are more accessible than the core features (Smith & Medin, 1981). The postulation of nonnecessary features answers many of the problems of the classical view. For example, the fuzziness of category boundaries could arise because of the presence of a nonnecessary feature of a particular exemplar of a category. Likewise, typicality/category goodness effects could arise for the same reason. However this additional assumption comes at a high cost with respect to the classical view: core features become less and less important thus rendering the classical view essentially irrelevant (Smith & Medin, 1981).

Perhaps the most influential work of this nature was that of Rosch and collaborators. Taking seriously the notion that the boundaries of categories are vaguely, rather than well, defined (Rosch, 1973), this work led to the notion of family resemblance. Their work was guided by the notion that categories are formed along two basic principles: cognitive economy (optimising the number of possible categories to a manageable extent) and real-world structure (the fact that many features of meaning naturally occur in tandem) (Rosch & Mervis, 1975). The best representatives of a category (prototypes) were found to be those exemplars which shared the most features of meaning with other members of the category, and shared the least with members of another category: a principle of family resemblance coupled with contrast to other non-family-members. This led to the notion of the basic level of representation: the level of specificity at which the combined within-category resemblance and between-category dissimilarity is the greatest (Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976). For example, in the hierarchy {*animal, mammal, pet, dog, collie, Lassie*}, *dog* would be the basic level. Most of the work by Rosch and colleagues focused on category membership and addressed some of the dominant theoretical controversies of the time, and many researchers of semantics have adopted assumptions and methodologies of this theory, particularly the possibility that category boundaries can be vaguely defined, and the resulting constructs related to family resemblance rather than all-or-nothing category membership.

A particularly useful application of Rosch's approach is the use of features of meaning as a tool to provide insight into word meaning. One method is to assemble sets of words from a semantic domain of interest and decide *a priori* upon their features but without claiming that these features constitute a complete set, and then use those features to build a computational model of representation which can then be tested against data. This method was taken by Hinton & Shallice (1991, also see Plaut & Shallice, 1993), who created a set of semantic features which intuitively capture properties of common objects (e.g. <has-legs>, <hard>, <made-of-metal>, <part-of-limb>), and used these features to train an attractor network to learn the mapping between orthography and semantics. Lesioning this network produced semantic, visual, and visual/semantic errors consistent with patterns of performance in deep dyslexia. Plaut (1995) used the same approach to investigate dissociations between reading concrete and abstract words. A particular characteristic of the representations was that abstract words had fewer features than concrete words. This difference in featural properties between concrete and abstract words (possibly in conjunction with other differences) translated into different consequences when different aspects of the model were damaged: abstract words were more impaired when the feedforward connections were lesioned, while concrete words were more impaired when the recurrent connections were lesioned. Such findings suggest that even double dissociations can arise from a model with a (single) common level of semantic representation, depending upon underlying characteristics of the featural input.

Theories based on independently-generated input. One possible problem with the computational models discussed in the previous section is the fact that the semantic features used were chosen by the investigators, and may, therefore, reflect the investigators' theoretical biases, or may not be true of the full range of meaning of the words in question. Other authors have addressed this concern by investigating those dimensions of meaning that are considered to be psychologically salient by others. Several models of semantic representations based on this kind of input have been implemented to date, differing mainly in the manner in which semantic representations are derived from the input. One class of models employs connectionist frameworks which develop representations from semantic input. These models are used in order to demonstrate how particular patterns of semantic impairment may be observed as a consequence of differential featural composition. This entails training a connectionist network with input that, although not directly obtained from speakers, is informed by characteristics of feature norms that are hypothesised to play a role. For example, Farah and McClelland (1991) constructed a model in which words referring to living or nonliving entities were associated with different proportions of visual-perceptual and functional features (the former predominant for living things, the latter predominant for nonliving entities). In this case, the proportions were derived from dictionary definitions, which were presented to naïve participants who were asked to rate the individual elements of meaning in each definition in terms of sensory/perceptual or functional content. When the model was lesioned, different category-related effects were found, depending upon whether the lesion targeted the visual-perceptual features (with living things more impaired) or functional features (with non-living things more impaired). A similar approach was taken by Devlin, Gonnerman, Andersen and Seidenberg (1998) who investigated the role of intercorrelated features (those features which frequently co-occur, e.g., <has wings> and <has a beak>) and distinguishing features (those which allow similar entities to be distinguished from each other) upon impairment over time for living or nonliving things as a consequence of dementia. Since living things have many intercorrelated features but few distinguishing ones, and the situation is reversed for nonliving entities (McRae, de Sa & Seidenberg, 1997), differences in their composition were able to explain the progression of relative impairment for living and nonliving things in dementia, within a single representational system (see also Rogers et al., 2004)

In these examples, semantic representations are based on specific characteristics derived from independently-generated information about meaning (e.g. more visual-perceptual features for living things, as in Farah & McClelland, 1991; more intercorrelated but fewer distinguishing features as in Devlin et al., 1998). Such approaches, however, require making *a priori* assumptions about the particular properties that are relevant to explain a particular pattern of data, and do not allow for the possibility that other properties not explicitly embedded in the semantic representations may also play important roles. Indeed, the theories of Farah & McClelland and Devlin et al. are not necessarily incompatible with each other, but their implementations do not permit direct comparison. This is because each model only embeds certain specific properties of featural input, and not others which are hypothesised to play a role under other theories, rather than simultaneously embedding multiple characteristics of featural input.

Another class of models based on independently-obtained input avoids the need of deciding in advance which properties are relevant to explain a given pattern of data by using speaker-generated features: separable aspects of meaning that naïve participants believe are important in defining and describing the meaning of a given word. These features are used to develop a model of representation, and then the properties of the resulting model are analysed to identify those properties that affect the representations (Hampton, 1979; 1981; Hampton & Gardiner, 1983; Rosch & Mervis, 1975; Rosch et al., 1976; Smith et al., 1974; Tversky & Hemenway, 1984). The first work along these lines to be carried out on a larger scale was conducted by McRae et al. (1997), who collected speaker-generated features for a large number of nouns referring to concrete objects (animals, plants, fruits, vegetables, artefacts, vehicles, etc.). Subsequent work by McRae and colleagues has used these features to address a number of questions of semantic representation and impairment (e.g. Cree & McRae, 2003; Cree, McNorgan & McRae 2006; Cree, McRae & McNorgan, 1999; McRae & Cree, 2002; McRae, Cree, Seidenberg & McNorgan, 2005; McRae, Cree, Westmacott & de Sa, 1999). Similar work based on speaker-generated features for nouns referring to objects has also been conducted by other groups (Garrard, Lambon Ralph, Hodges & Patterson, 2001; Randall, Moss, Rodd, Greer & Tyler, 2004; Rogers & McClelland, 2004), collectively providing comprehensive data sets which allow investigation of semantic representations from numerous directions. However, one crucial characteristic of many of these models is that they still tend to discuss the featural input in terms of one particular dimension (or just a few) to explain particular patterns of data. To permit the evaluation of such models more generally as theories of semantic representation, it is necessary to consider classes of models which do not depend on selecting particular characteristics of the featural input, but which still permit features to be analysed in such terms if desired (Vinson & Vigliocco, 2002; 2008; Vigliocco et al., 2004). This is one of the central aims of the present work.

Going beyond semantic representations for words referring to objects

Nearly all of the theoretical and behavioural research described above has focused upon nouns referring to concrete objects. As Miller and Fellbaum (1991) put it, "When psychologists think about the organisation of lexical memory it is nearly always the organisation of nouns they have in mind." (p.204). But nouns referring to concrete objects are only one semantic domain, and it is unclear whether theoretical conclusions based only upon investigations of the semantics of concrete nouns can be generalised to other semantic domains. Words referring to actions[3] is the only other domain beyond words referring to concrete objects that has received some attention (although see Gross, Fischer & Miller, 1989, for some discussions on the domain of words referring to properties). A first difference between words referring to objects and words referring to actions is referential: words referring to objects can be understood in isolation, while words referring to actions are relational in nature. One implication of this difference is that words referring to actions are more abstract than words referring to objects (Bird, Lambon Ralph, Patterson & Hodges, 2000; Breedin, Saffran & Coslett, 1994). Some authors have also argued that words referring to objects and actions differ in featural properties (Graesser, Hopkinson & Schmid, 1987; Huttenlocher & Lui, 1979). For words referring to objects there are more features referring to narrow semantic fields (e.g., <domesticated> vs. <wild> for animals). For words referring to actions, instead, there are more features that broadly apply across a wide range of semantic domains (e.g., <intentionality>, <involves motion>). As a consequence of this, the patterns of correlation among semantic features would also differ for words referring to objects and actions; features should be much more strongly correlated within semantic fields for words referring to objects (e.g., <having a tail>, and <having four legs> for mammals) than for words referring to actions.

Research on categories also demonstrates that distinguishing between different levels (superordinate, basic, subordinate) (Rosch & Mervis, 1975; Rosch et al., 1976) is relatively simple for words referring to (most) concrete objects; the situation is different for words referring to actions, for which it is typically difficult to create comparable sets of hierarchies. For example, *yell* does not fall easily into a superordinate category, as it can be considered *communication*, *noise*, or *mouth action*, any one of which seems insufficient as a category label compared to something like *animal* or *fruit*. The word *yell* also appears to lack subordinates (more specific instances of yelling would probably be reflected in the use of modifiers rather than in selection of a different word). Nonetheless, there have been attempts to define basic level actions (Lakoff, 1987; Morris & Murphy, 1990) and some attempts have been made to describe words referring to actions in hierarchical terms (Jackendoff, 1990; Keil, 1989). Differences between the domains, however, persist. For example, in Keil (1989), the hierarchical organisation of words referring to actions has fewer levels (generally two) and with fewer distinctions at the superordinate level. Other attempts to capture a level of organisation for actions have included distinctions between light (e.g., *do*) and heavy (e.g., *construct*) verbs (Jespersen, 1965; Pinker, 1989) and distinctions between general

[3] The term "action" is used here in a broad sense, encompassing both words referring to physical actions (e.g. *run, throw, eat*) and events of various kinds (e.g. *clatter, glow, preach*).

(e.g., *move*) and specific (e.g., *run*) actions (Breedin, Saffran & Schwartz, 1998). However, the light/heavy dichotomy only allows drawing a distinction between verbs used as auxiliaries and those that participate in phrasal verbs (e.g., *get up*, *throw away*) on one hand and all other verbs on the other (and permits no such distinction for nouns referring to actions), and drawing the line between general and specific actions is not an easy or agreed-upon exercise. A related issue is that distinctions between close semantic neighbours differ across the domains of words referring to objects and to actions. For many domains of basic-level concrete objects, close neighbours offer true distinctions (a goat is not a sheep; an apple is not a pear) while this is not true in many action domains which seem to overlap to a greater extent (e.g. *to shout, to yell, to scream*; none of which necessarily exclude any of the others).

Words referring to objects and actions also differ in syntactic terms, particularly considering that all object words are nouns, while the typical action word is a verb. As such the syntactic information associated with words referring to actions tends to be richer than for words referring to objects: the lexical-semantic representations of actions are considered to contain not only the core meaning (the action or the process denoted) but also the thematic roles associated with the verb. For example, the core meaning of *to kick* is something like "striking out with the foot" and it is associated with the thematic roles of Agent and Patient, the arguments of the verb that specify "who did what to whom". (Grimshaw, 1991; Jackendoff, 1990; Levin, 1993). The same could also be said of nouns referring to actions, which despite fulfilling the same syntactic roles as nouns referring to objects (e.g., subject, object, head of noun phrases) also take arguments in the same manner as verbs referring to actions (Collina, Marangolo & Tabossi, 2001).

Such differences between the semantic representations of words referring to objects and words referring to actions have led to some independently-developed and distinct models of representation for the two. As discussed above, in Wordnet (Miller & Fellbaum, 1991), different types of relational links have been implemented for objects and actions to accommodate the various differences between the domains. For nouns referring to objects, relations such as synonymy, hyponymy (e.g., *dog* is a hyponym of *animal*), and meronymy (e.g., *mouth* is part of *face*) are argued to play a primary role in semantic organisation. For nouns and verbs referring to actions, relational links include troponymy (i.e., hierarchical relation in which a subordinate term expresses the manner of its superordinate, such as the relation between *crawling* and *travelling/ going/ moving/ locomoting*), entailment (e.g., *snoring* entails *sleeping*) and antonymy (e.g., *coming* is the opposite of *going*). It is possible, however, that an assumption of separate representational systems is not necessary. After all, work by researchers using speaker-generated features has revealed that substantially different semantic domains (e.g. living vs. nonliving things) can be represented in a single model, using a common set of implementational assumptions. Differences between these domains of knowledge come about because of differences in the types of properties speakers generate for words in a given domain. The present work extends this notion even further, investigating whether even more diverse domains of meaning, words referring to objects and words referring to actions, can be successfully represented in a single system.

The present work

Moving on from the extensive work by McRae and colleagues, the central question in the present study is whether the same speaker-generated semantic feature approach is a suitable way to investigate and model semantic representations for words referring to objects and actions despite the numerous differences between the two domains. This question will be addressed by a series of parallel studies based on speaker-generated features, collecting and analysing features of words referring to objects (for which much is known thanks to the work by McRae and others), and also words referring to events, in order to assess differences between the domains, and to find out if the representational assumptions provide similarity measures of comparable quality for object-nouns and for words referring to actions. Chapter 2 introduces the item set and describes the feature collection methodology. Chapter 3 contains analyses of the nature and content of features that were generated for words from various semantic domains. In Chapter 4, an implementation of a model with distinct conceptual and lexical-semantic levels of representation is introduced and described. Chapter 5 explores the characteristics of the resulting lexical-semantic similarity space at different levels of specificity. Chapter 6 presents four behavioural experiments that tested the ability of the model to predict fine-grained word-level semantic similarity effects in comprehension and production, and compares the performance of the model for words referring to objects and words referring to actions. Chapter 7 uses data from the experiments in Chapter 6 to assess the quality of the speaker-generated feature model against two other models from which word-level semantic similarity measures are available (Latent Semantic Analysis (LSA), and Wordnet). Chapter 8 presents an additional experiment exploring the ability of the speaker-generated feature model to predict category-level semantic effects, and Chapter 9 offers discussion and conclusions.

Chapter 2: Feature collection

In order to provide a suitable basis for a model of semantic representation based on speaker-generated features, it is first necessary to select the model's vocabulary from the entire English lexicon. After all, collecting and processing speaker-generated features is a very time-consuming process. At the same time, however, it is necessary that the words included in a model are suitably broad in scope and relevant to important issues in the literature.

We begin with nouns referring to concrete objects, as a substantial literature already exists concerning this domain, particularly concerning category-related deficits, and because several sets of speaker-generated feature norms already exist to serve as comparison (e.g., Garrard et al., 2001; McRae et al., 1997; Randall et al., 2004; Rogers & McClelland, 2004). Patterns of category-related deficits suggest that it is especially important to include a variety of living and nonliving things, including fruits and vegetables, animals, clothing, furniture, vehicles, tools, and other artefacts. The field of body parts is particularly interesting in this regard, belonging to living things but for which it might be argued that functional properties are particularly important, thus more like nonliving things according to the distinctions described by Farah and McClelland (1991). Other considerations for selection of words referring to objects were that they should be familiar (so that participants would be able to generate features for the words rather than generic features referring to superordinates), relatively unambiguous or at least with a dominant meaning (so that participants would not produce diverse sets of features referring to different meanings of a given word), and ideally picturable (to permit their use in behavioural experimentation). A variety of exemplars were included within each category in order that a range of semantic similarity would be represented within the set, thus allowing the quality of the resulting model to be assessed at a variety of levels of specificity.

Since less is known about the semantic organisation of words referring to actions, it was more difficult to decide which of these to include in the list. This selection process began with verbs referring to actions. An initial set were selected because they describe actions that are associated with words already included in the object set (e.g., words related to cooking, to the use of tools, and body actions). Other words referring to actions were selected because intuitively they offered variability in their featural composition: words referring to light and sound emission which were expected to have sensory-related features; words referring to manner and direction of motion for which motion features are expected to be important; and words referring to communication and exchange for which features related to purpose/function may dominate. Again, words were selected with familiarity and limited polysemy in mind, and picturable actions were selected where possible. Finally, a set of nouns referring to actions were included in the set. All of these were homonymous or derivationally related to the verbs in the list (e.g., *plea/plead*; *donation/donate*).

Method

Item selection

A total of 456 words were selected. This list included 216 verbs referring to actions and 240 nouns, 169 referring to objects and 71 to actions. A complete list of the words, along with their semantic field labels, is given in Appendix A. The 456 words were pseudorandomly assigned to 14 lists, each of which contained 30 to 40 words. Nouns referring to objects, nouns referring to actions and verbs referring to actions, as well as exemplars from any given semantic field[4] (i.e., animals, tools, verbs of body action, verbs of light emission etc.) were distributed across the lists as evenly as possible. In English, many noun and verb forms are homonymous, therefore, in order

[4] Henceforth the term "semantic field" will be used to refer to groups of words related in meaning. For many types of object nouns this term is synonymous with "category", but this latter term does not apply so clearly when actions and events are also considered.

to obtain features for both the noun and verb versions of the same word form, all nouns and verbs were disambiguated by presenting them in a minimal syntactic context: verbs were presented in the infinitival, with *to*, and nouns were preceded with the determiner *the*. Noun-verb homonyms (e.g., *the hammer*, *to hammer*) were always assigned to different lists, as were derivationally-related nouns and verbs (e.g. *donate/donation*). The order of items in each list was randomised, and words were printed six to a page, each with ten blanks in which participants were to record their features.

Procedure

Two hundred eighty undergraduate students from the Department of Psychology at University of Wisconsin, Madison, participated in exchange for extra credit. Twenty participants completed each list. For each item they were asked to write down those features of meaning which, taken in conjunction, were sufficient to define and describe that word. The instructions defined features as "words or phrases that, taken alone, provide a single piece of meaning information". Features for two examples (one noun referring to an object and one verb referring to an action, neither occurring on that list) were provided as a model[5]. Participants were instructed to avoid producing pure associations ("for the word 'cat', you would not produce the feature 'mouse', because although cat and mouse are related, 'mouse' does not reflect what a cat is"). Although the response sheets contained ten blank lines for each word, participants were not explicitly instructed to produce ten features per word, but rather, to produce "enough features to define and describe the word" (with a maximum of ten responses). Finally, participants were instructed to produce features for each word in turn, and not to return to a word once they had begun generating features for the following word. The task lasted approximately 45 minutes. Five participants who failed to comprehend the task were replaced.

[5] These features were obtained on the basis of features generated for items during a pilot study.

When participants missed a word, generated features for the wrong word (or wrong meaning), or did not know a word, that word was added to one of three follow-up lists. Follow-up lists were presented to additional participants to ensure that there were 20 participants' responses for each word in the set. During the feature collection process some words were shown to be problematic. For example, *to trill* was not known by almost half of the participants, and features for *the twist* almost always referred to a type of dance rather than to the intended noun depicting the physical action of twisting (features for the related verb *to twist* never referred to the dance). Features for *the foot* were (almost) evenly divided between referring to the body part and to the unit of measurement. Words with problems of this type were eliminated from the set. Although an effort was made to select words that were not polysemous, some degree of polysemy was unavoidable. In order to minimise the effects of polysemy, participants were instructed with the following, "Sometimes you might think of more than one meaning of a particular word. When this occurs, please write down features only for the most common meaning of this word, in your opinion". Indeed, nearly all participants generated features only for the most common meaning. Words for which some participants generated other meanings were added to one of the follow-up lists.

Analysis

Data consisted of a large quantity of hand-written speaker-generated features for each word from 20 participants. Initial data entry consisted of entering each feature into a feature x participant matrix for each word. Thus, this phase of data entry reflected the participants' intent as closely as possible. When a participant produced a given feature for a given word, a value of 1 was entered for that feature x participant cell (indicating that the feature was present), otherwise that cell was left blank and subsequently a value of 0 was entered (indicating that the feature was absent).

Once the numeric data were entered for each word, the post-processing of the featural data involved making a number of decisions in order to capture aspects of similarity among speaker-generated features. For example, individual participants occasionally produced conjoint features (i.e. <red fruit> for *the apple*), while others produced such features separately (<fruit>, <red>, etc.). If other participants produced both features separately for the same word (as in <red fruit> above) these cases were considered to be unambiguously separable, and were converted into their separate features. For those situations in which this was not the case (e.g., if one participant produced <red fruit> but no other participant produced both <red> and <fruit> as separate features) this was decided on a case-by-case basis: a decision of whether the conjoint features (<red fruit>) had a substantially different meaning than its contributing features (<red> and <fruit>) taken separately. When a conjoint feature was separated, a value of 1 was entered for both features for that participant. Along similar lines, participants produced a wide range of variations in wording or synonyms to express the same feature (e.g. <4-legged>, <has four legs>, <quadruped>, etc.). Synonymous features were collapsed into a single feature, and the shortest, most-frequent variation was chosen from the different alternatives. This procedure was carried out for each of the 456 words in the set. Sometimes, a participant produced two (or more) synonymous features in response to a single target word. Such cases were treated as if the speaker had produced the feature only once (in other words, a binary coding of present or absent for each participant).

Because the above analyses were carried out on a word-by-word basis, the treatment of synonymous features proved somewhat problematic: a "synonymous feature pair" for one word might be missed when another word is considered (for example, if <fast> and <quick> are considered synonymous, and all instances of <quick> converted to <fast>, this could easily be missed for a word for which the feature <quick> is produced but <fast> is not). Therefore, once all synonymous features were dealt with, a complete list of features was prepared, consisting of one instance of each feature that occurred across the entire set of words. This list was examined independently from the target words, and each feature was examined in relation to all other features in the list, to identify further possible synonyms, on an intuitive basis and with the aid of a thesaurus (for example, given the feature <fast>, the list was searched for other words referring to speed such as <quick>, <speedy>). A feature was considered as "synonymous" if it appeared in a thesaurus entry for another word such as <noisy> for <loud>. When a feature pair was judged to be synonymous, it was changed across all target words in the set. Feature weight values were determined for each feature in that word, as the number of speakers who generated that feature for that word (for example, 19 participants produced the feature <red> for *the cherry*, so the feature <red> was given a weight value of 19 for cherry). Feature weight vectors for each word were then prepared by enumerating all of the feature weight values of that word, assigning values of zero to all features not produced for a given word. Finally, a word x feature matrix was created by combining the feature weight vectors across all 456 words in the set. In this matrix, values represented the number of speakers who had produced a given feature for a given word. At this point, idiosyncratic features (those produced by nine or fewer participants across all words, i.e., with summed feature weight values of less than 9) were discarded, resulting in a feature weight matrix of 456 (words) by 1029 (features). Features and their weights for each word are published as online supplementary materials accompanying Vinson and Vigliocco (2008; see Appendix A for details).

Chapter 3: Properties of the speaker-generated features.

Once the feature x word matrix was completed, it was then possible to evaluate the extent to which properties of the features differ across words from different domains of meaning (words referring to objects and actions), and across words from different semantic fields (e.g., words referring to animals, body parts, artefacts; manner of motion, tool action, communication, etc.). Beyond providing descriptive information about the different domains and semantic fields, analysis of the featural space also allows us to assess specific claims within cognitive psychology and neuroscience. These include differences between semantic representations of objects and actions (both in terms of their content and their organisation) and between living and nonliving domains for which many specific claims have been made. They also allow us to assess claims concerning category-related deficits, such as the different role of sensory, functional and motoric features, and different patterns of correlations among features.

Here we report the following analyses:

Characteristics of the general featural properties of the words themselves (number of features and total weight of features). These general properties can reveal aspects of representation such as semantic richness which might presumably differ across domains of knowledge. For example, words referring to objects might be semantically richer than words referring to actions because the latter are relational (and thus not containing semantic content related to the participants involved in the activity); this could be reflected in higher number/weight of features for object-nouns than for action-words. The reverse could also be true, if words referring to actions also include content related to the participants involved in the activity. These analyses can also reveal fine-grained differences between semantic fields in a broad domain (e.g., whether different categories of words referring to objects exhibit different featural characteristics, or if categories of words referring to objects are homogeneous in this regard).

Distribution of features of a given type (visual, other perceptual, motional and functional) among words in a given field. A number of theories explaining differential patterns of impairment following brain injury are based upon distinctions between different feature types (e.g. the sensory-functional theory, Warrington & Shallice, 1984) Analysis of these characteristics of words across domains of words referring to objects and to actions, and also fine-grained analysis at the level of semantic fields, can provide further insight into the applicability of these theories which have mainly addressed representations of nouns referring to objects only.

Relationships among features in words: the extent to which words of different domains share features, the extent to which features correlate, and the role of distinctive features in words from different domains. Some theories rely upon these aspects of meaning to explain differential patterns of impairment, though only in the domain of nouns referring to objects, as discussed below.

General featural properties

Number of features. First the number of features generated for each word regardless of the number of participants who generated them was calculated. The average number of features per word was 28.33 (SD=6.09). Second, the average number of features for individual words in a given semantic field was calculated, as a measure of semantic richness. The first column of Table 1 gives a summary of mean number of features, organised by semantic field.

Table 1. Average number of features, average sum of feature weights, and their ratio, as a function of semantic field (standard deviation in brackets)

	Mean # features	Mean feature weight	Weight/number ratio
Objects			
Animals	30.96 (7.47)	126.2 (24.65)	4.08 (0.70)
Fruit & vegetables	26.88 (5.00)	115.8 (21.59)	4.31 (0.69)
Tools	30.48 (6.44)	109.9 (18.56)	3.61 (0.60)
Vehicles	31.00 (6.41)	109.8 (24.62)	3.54 (0.53)
Body parts	32.46 (5.93)	116.4 (18.40)	3.59 (0.65)
Clothing	27.18 (5.21)	108.4 (16.69)	3.99 (0.49)
Misc. artefacts	32.45 (5.17)	106.6 (17.73)	3.29 (0.59)
Actions			
Body actions	29.15 (5.80)	98.0 (17.23)	3.36 (0.86)
Body sense	26.17 (7.38)	89.7 (20.01)	3.43 (0.90)
Change of location	29.36 (7.29)	86.4 (15.39)	2.94 (0.68)
Change of state	25.40 (4.86)	76.4 (15.93)	3.01 (0.36)
Noises	28.17 (3.58)	93.6 (29.10)	3.32 (0.91)
Communication	28.89 (7.97)	88.8 (23.36)	3.07 (0.74)
Construction	27.29 (6.00)	95.1 (26.23)	3.48 (1.00)
Contact	27.33 (5.92)	91.5 (17.77)	3.35 (0.84)
Cooking	24.43 (5.79)	95.9 (18.25)	3.93 (0.80)
Destruction	31.88 (6.52)	89.4 (24.61)	2.80 (0.59)
Exchange	23.50 (5.39)	80.3 (16.92)	3.42 (1.00)
Heat/light emission	25.54 (6.09)	88.8 (20.04)	3.48 (0.48)
Motion direction	22.00 (5.32)	73.6 (25.83)	3.35 (1.02)
Motion manner	29.12 (6.78)	95.9 (23.33)	3.29 (1.06)
Tool action	34.22 (5.59)	104.1 (14.03)	3.04 (0.51)

In order to statistically compare the number of features for words from different semantic fields, we started by formally comparing object and action words using a two-tailed t-test. To examine finer-grained semantic field effects, we conducted separate analyses for objects and actions, first conducting omnibus F tests comparing all of the semantic fields listed in Table 1. When an omnibus F-test was significant, we followed up with one-tailed t-tests (Bonferroni-corrected) contrasting semantic fields that appear to exhibit large numeric differences on a given measure. This approach is necessary in order to reduce the number of comparisons to a manageable level; there are 105 possible comparisons involving pairs of the 15 action semantic fields listed, and 21 possible comparisons involving pairs of object fields. Comparisons were therefore restricted to testing apparent differences, and cases for which specific claims have been made in the literature. This same procedure was also carried out for the other semantic field comparisons reported later in this chapter.

In general, more features were generated in response to words referring to objects than to words referring to actions ($t(383) = 3.4196$, $p < .001$), though this was not true for all semantic fields. For example, speakers generated numerically more features for tool actions than for any of the object fields, and fruit/vegetables and clothing had an average number of features similar to that of an average action field. Among words referring to objects, the highest number of features were produced for body parts and miscellaneous artefacts, and the fewest, for clothing and fruit/vegetables (Body vs clothing $t(38) = 2.59$, $p = .014$; body vs. fruit/veg $t(57) = 4.64$, $p < .001$; misc artefacts vs clothing $t(36) = 2.50$, $p = .017$; misc artefact vs. fruit/veg $t(55) = 4.41$, $p < .001$). Among the words referring to actions, the highest number of features were produced for tool actions and actions referring to destruction, and the fewest, for more abstract words, such as those from semantic fields like direction of motion and exchange (tool action vs. motion direction $t(22) = 6.08$, $p < .001$; tool action vs exchange $t(23) = 4.54$, $p < .001$; destruction vs motion direction $t(21) = 4.60$, $p < .001$; destruction vs exchange $t(22) = 3.33$, $p = .003$).

<u>Summed feature weight</u>. Although feature number indicates semantic richness to some extent, it does not capture the whole nature of semantic representations, because the binary distinction between presence/absence of a feature for a word does not take into account how salient features are (see Smith & Medin, 1981 for a number of arguments in favour of variable, rather than binary, values of semantic features). This is especially important because for most words in the set, the vast majority of features had very low weights (i.e., they were produced only by a few participants). For example, the word *lemon* had a total of 31 features, but only three of them had weights greater

than 10 (<yellow>, <sour>, <fruit>) and only three more had weights greater than 5 (maximum = 20). Here, feature weights (the number of participants who produced a feature for a given word) were used as a more informative measure of semantic composition and a more precise reflection of the underlying meaning representations of words. The measure used here is summed feature weights: the total number of different features produced by all participants for a given word (after the post-processing procedure and removal of idiosyncratic features described in chapter 2).

Unlike the number of features, feature weights clearly distinguished between words referring to objects and words referring to actions. All object fields exceeded all action fields in mean summed feature weights and they were significantly different when analysed by items ($t(383) = 13.32$, $p < .001$).(although this was not always true of individual exemplars of low-weighted objects such as *ceiling* [70], *shield* [73], *wing* [85], *tail* [85], *dress* [87] and a few high-weighted actions such as *breathe* [140], *speak* [132], *swim* [129], *write* [127], *cook* [124]). Again, fine-grained semantic field distinctions were observed, but not necessarily in the same way as for number of features. Within the object domain, largest weights were observed for animals, fruit/vegetables and body parts, and lower weights for tools, vehicles, clothing and other artefacts (Animals were significantly higher than all four of the latter categories (vs tools $t(48) = 3.44$, $p = .001$; vs vehicles $t(36) = 2.90$, $p = .006$; vs clothing $t(39) = 3.33$, $p = .002$; vs misc artefacts $t(45) = 4.03$, $p < .001$). Fruit/veg were significantly greater than misc artefacts only: $t(55) = 2.33$, $p = .024$, all other $p > .10$. Body parts exhibited the same pattern: vs. misc artefacts $t(44) = 2.49$, $p = .017$, all other $p > .05$). For actions, largest weights were observed for tool actions, body actions, cooking, manner of motion, and lower weights for change of state, direction of motion and exchange (Tool action vs change state $t(17) = 4.84$, $p < .001$; tool action vs. motion direction $t(22) = 6.28$, $p < .001$; tool action vs. exchange $t(23) = 3.73$, $p = .001$. Body action vs change state $t(48) = 3.24$, $p = .002$;

body action vs. motion direction $t(53) = 4.45$, $p < .001$; body action vs exchange $t(54) = 3.13$, $p = .003$; cooking vs change of state $t(15) = 3.30$, $p = .005$; cooking vs motion direction $t(20) = 4.41$, $p < .001$; cooking vs exchange: $t(21) = 2.24$, $p = .035$ (n.s. after correction for multiple comparisons); manner vs change state $t(33) = 2.94$, $p = .006$; manner vs direction $t(38) = 4.08$, $p < .001$; manner vs exchange $t(39) = 2.66$, $p = .011$) . The second column of Table 1 above gives the average feature weights of the different semantic fields.

The relative difference between the values of feature numbers and feature weight (for example, fruits/vegetables had the lowest number of features produced yet were among the highest in terms of summed feature weight), is due to the fact that feature weight takes into consideration not only the number of features generated but also inter-participant agreement. Some words elicited relatively small number of features overall, but for which participants were largely in agreement. This was reflected in the relatively high weights assigned to those features (e.g., *zebra* had only 19 features but with a high summed weight of 147; *shirt* had only 17 features with weight of 113; *peach*, 21 features with weight of 136). Words with high levels of inter-participant agreement tend to be uniquely defined and concrete nouns referring to objects. Other words may have elicited many features, but they were not highly weighted, because of less agreement among participants as to their meaning characteristics (e.g. *argue* elicited 33 features with a summed weight of only 57; *preach*, 36 features with weight of 66; *give*, 35 features with weight of 74). Because of the mismatch between feature number and feature weight, the ratio of weight per feature was also calculated for each word, as illustrated in the last column of Table 1. The largest ratio of weight to features was observed for words referring to objects (objects vs actions: $t(383) = 7.93$, $p < .001$): living things (animals and fruit/vegetables) with a ratio above 4 and other object fields with a ratio above 3.5 (except for miscellaneous artefacts, with a ratio of 3.29). In contrast, no action field exceeded a ratio of 3.5, and much more variability was observed among semantic fields. Construction and heat/light emission had the highest ratio (3.48), and destruction, change of location, change of state and tool actions had the lowest (all below 3.1) (Construction vs destruction $t(13) = 2.77$, $p = .016$; construction vs change location $t(16) = 2.62$, $p = .018$; construction vs change state (n.s.) $t(15) = 1.96$, $p = .069$; construction vs tool actions (n.s.) $t(14) = 1.88$, $p = .080$; heat/light emission did not significantly differ from the others due to the small number of words in this category).

Taken together, these results indicate that most words referring to objects in the set - measured in terms of the speaker-generated features - are semantically richer than words referring to actions: they had more features, with greater weights, and a higher ratio of weight to features. This is consistent with previous claims in the literature concerning differences between object and action representations (e.g. Plaut, 1995), and is the first indication that some of the important differences between words referring to objects and words referring to actions are reflected in basic properties of speaker-generated features.

Types of features

A second question of interest was whether words from broadly different domains (objects and actions) and within different semantic fields (animals, tools, communication, manner of motion, etc.) differ in terms of feature types. Previous studies (discussed in more detail below) suggest that three major differences are expected: (a) more perceptual features are expected for words referring to objects than for words referring to actions, (b) more perceptual features are expected for living things than for artefacts, and (c) motion features are expected to be more common among the words referring to actions than among the words referring to objects.

Identifying the feature types associated with different domains and semantic fields would allow the evaluation of the claims of the sensory-functional hypothesis that has been put forward to account for category-specific deficits (Warrington & Shallice, 1984). In order to do so, a finer-grained criteria than in previous studies (e.g., Farah & McClelland, 1991; Garrard et al., 2001) was used for classifying the speaker-generated features. Instead of distinguishing only between sensory and functional features, the present study made further distinctions between functional and motoric features, and also between visual and other perceptual features. The contrast between motoric and functional features was introduced because of evidence that knowledge of how to use an object (motoric) and knowledge of what the object is used for (functional) can dissociate in some patients (Buxbaum, Veramonti & Schwartz, 2000), an especially important finding considering that motoric features are not mentioned by sensory-functional accounts of impairment. Note that to some extent, the distinction between functional and motoric features may correspond to explicit and implicit knowledge of action, and both functional and motoric features tend to be very general in nature (for example, the specificity of motoric features related to grasping are limited to the observation that a particular implement is used with the hand, rather than any further details such as hand configuration, orientation, muscles used for action, and so on). The speaker-generated features were classified into five categories by two native English speakers. Any disagreements were discussed and agreed upon. First, all perceptual features, using the definition of "features that describe information gained through sensory input, including body state and proprioception" were identified. Perceptual features were subdivided into Visual Features (constituted 22.2% of all features), and Other Perceptual Features that refer to any other sensory modalities (which constituted

19.7% of all features)[6]. Next, features were classified into Functional (features referring to the purpose of a thing, "what it is used for", or the purpose or goal of an action. These constituted 26.5% of all features), Motoric ("how a thing is used, or how it moves", or any feature describing the motor component of an action. These constituted 12.0% of all features), and Other Features (those meeting none of the previous classification schemes, constituting 37.6% of all features) The class of Other features contains the largest proportion of all the features, and is highly heterogeneous. Some of the features classified as Other are encyclopaedic (e.g., [comes from] <Africa>); while others refer to relationships among meaning components, (e.g., ISA <animal>; PART OF <face>, relationships that are particularly common in taxonomies developed by lexicographers; (see Miller & Fellbaum, 1991). As such these are highly variable among items.

For the purpose of the current work, Other Features were not further classified, since these do not play a role in previous theories of semantic organisation. Figure 1 represents the distribution of feature types in object semantic fields, and Figure 2 represents the distribution of feature types in action fields, (taking weights into account). As can be seen in the following Figures (see next page), words referring to objects and to actions appear to differ in their featural composition.

[6] These feature type classifications were not mutually exclusive; for example, some sensory properties can be experienced through multiple channels, such as <smooth> which has visual implications as well as tactile. Features of this kind were permitted multiple classifications, so the total number of features exceeds 100%.

Figure 1. Percentage of feature types in exemplars from various object semantic fields, adjusted by weight. Error bars reflect standard error of the mean by items.

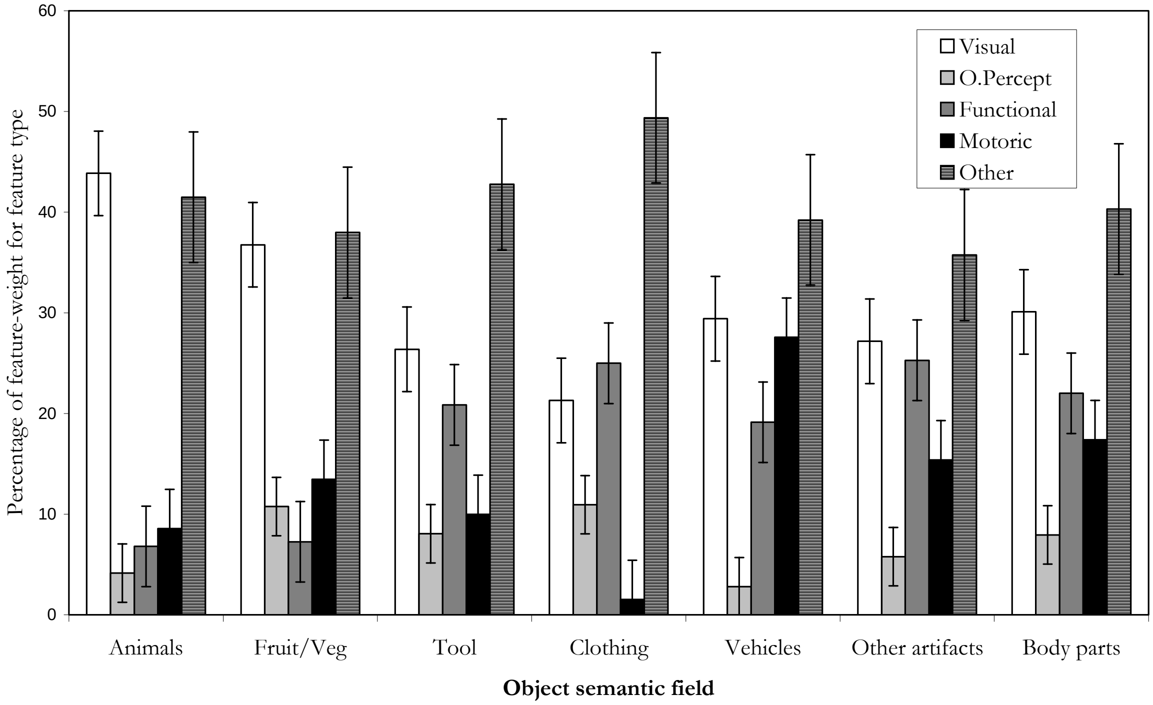

Figure 2. Percentage of feature types in exemplars from a subset of action semantic fields, adjusted by weight. Fields were selected to be indicative of the range of featural composition in the complete set of semantic fields. Error bars reflect standard error of the mean by items.

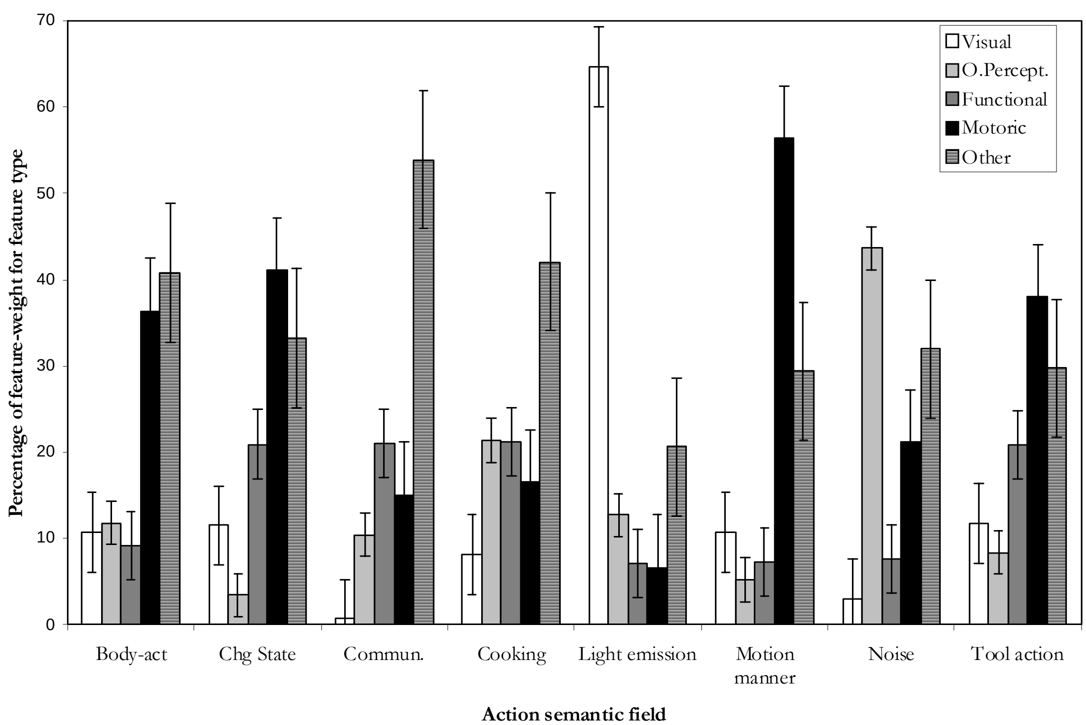

Perceptual features. First, all perceptual features were considered together. Words referring to objects were found to be more dependent upon sensory features than words referring to actions ($t(454) = 6.215$, $p < .001$)[7], with the exceptions of light emission (77.3% of all features) and noise making words (46.6%) which were the only action categories above the group mean for objects. Within object domains, living things were most dependent upon sensory features (48.0% for animals, 47.5% for fruit and vegetables), and other artefact fields appeared to be less so (tools 34.5%, clothing 32.2%, vehicles 32.2%, other artefacts 33.0%), while body parts had an intermediate number of perceptual features (38.0%). However none of these within-category differences reached significance, likely due to variation of individual items within each class (all p-values > .03, not significant when corrected for multiple comparisons). The difference between living and nonliving things (excluding body parts) was significant; $t(142) = 5.860$; $p<.001$. This finding is consistent with the claims of the sensory-functional hypothesis. In contrast to the object fields, most action fields had substantially lower proportions of sensory features (e.g., change of state, 14.9%; communication, 11.1%), suggesting the prediction that action naming (relative to object naming) should be relatively spared when sensory features are impaired. Nevertheless, despite the finding in the present study that some words referring to actions (e.g., noise making, communication, light emission labels) had far more sensory features than any object domain, there are no reports in the literature of impairments of words belonging to these semantic fields associated with impairments to sensory features (possibly because no one has tested patients on words from these fields). Next, more specific analyses were conducted, focusing on visual features (the dominant modality among the perceptual feature types).

[7] All t-tests reported in this thesis are two-tailed unless otherwise specified.

Visual features. Among the words referring to objects, visual features were most salient for animals (43.9% of all features) and fruits and vegetables (36.8%), moderately salient for vehicles (29.4%) and body parts (30.1%), and less so for other (artefact) fields (26.4% for tools, 21.3% for clothing and 27.1% for miscellaneous artefacts), (animals vs. clothing: $t(39) = 4.210$, $p < .001$; animals vs tools: $t(48) = 3.671$, $p < .001$; animals vs misc artefacts $t(45) = 3.815$, $p < .001$; fruit/veg vs clothing $t(49) = 2.991$, $p = .004$; fruit/veg vs tools $t(58) = 2.656$, $p = .010$; fruit/veg vs misc artefacts $t(55) = 2.42$, $p = .019$)This tendency is in line with the distinction between living and nonliving categories and with the claims of the sensory-functional account. Living things had significantly higher weighted feature composition than nonliving things, $t(142) = 4.151$, $p < .001$ (this comparison includes all object fields except body parts). Fine-grained differences were also observed *within* these fields: among living things, animals were more dependent upon visual features than fruits and vegetables ($t(57) = 2.990$, $p = .002$) and among artefacts, vehicles were more dependent upon visual features than tools or clothing ($t(57) = 2.451$, $p < .001$). Considering words referring to actions, visual features were predominantly salient only for the narrow semantic field of light emission (e.g. "glow", "shine"), for which visual features were by far the dominant feature type, amounting to 64.6% of all features. Other action fields had little, if any, dependence upon visual features. This finding lends credence to the suggestion that loss of visual features can result in category-specific impairments of object naming, and within object fields, of living things (e.g. Allport, 1985; Farah & McClelland, 1991).

Other perceptual features. In relation to other perceptual features, among words referring to objects, again, consistent differences between fields were observed. Fruit and vegetable (10.8% of all features) and clothing (10.9%) were most dependent upon nonvisual perceptual features, tools (8.1%) and body parts (7.9%) moderately so, but other fields less so (animals, 4.2%; vehicles, 2.8%; other artefacts, 5.8%). Again, some fine-grained differences were observed within artefact and living domains: e.g., clothing had significantly more other-perceptual features than tools ($t(44) = 1.813$, $p = .038$); tools had more than vehicles ($t(41) = 2.162$, $p = .019$); fruit/vegetable had more than animals ($t(57) = 4.684$, $p < .001$). Several action fields, noise making (43.6%), cooking actions (21.4%), communication (10.4%) and light emission (12.7%), were more dependent than any object field on nonvisual perceptual modalities.

Functional features. Artefacts and living things were consistently distinguished insofar as nonliving things were more dependent on functional features than living things ($t(142) = 8.152$, $p<.001$). Among the nonliving things, clothing (25.0% of all features) and miscellaneous artefacts (27.2%) were most reliant upon functional features, followed by tools (20.8%), body parts (22.0%) and vehicles (19.1%). Clothing had significantly more functional features than either tools or vehicles (respectively, $t(44) = 1.982$, $p = .031$; $t(27) = 1.996$, $p = .028$). Living things - animals (6.8%) and fruit & vegetables (7.3%) were considerably less dependent on functional features, significantly differing from all of the nonliving categories (all pairwise comparisons yielded $t > 3.5$, $p < .001$, significant after correcting for multiple comparisons). Considering words referring to actions, purposeful acts such as change of state (20.9%), communication (21.0%), cooking (21.2%) and tool actions (20.9%) relied most heavily on functional features. Other action fields relied on functional features to a considerably less extent (e.g., light emission (7.1%), noises (7.6%) and manner of motion (7.3%), all of which significantly differed from change of state, communication, cooking and tool actions; all pairwise comparisons yielded $t > 3$, $p < .001$). Overall, functional features were more important for the semantic makeup of words referring to objects than of words referring to actions ($t(454) = 2.774$, $p = .003$). This is contrary to the suggestion of Bird, Howard & Franklin (2000) that functional features are similarly important for actions and for inanimate objects (and that the loss of functional features should result in impaired performance for artefacts and actions alike). Functional features were also more important for artefacts than for living things, a finding consistent with the suggestion that loss of functional features selectively affects artefacts (e.g. Allport, 1985; Farah & McClelland, 1991). Again, however, reports of impairments to action fields such as change of state, communication, cooking and tool actions, all of which have a relatively large proportion of functional features (in numbers comparable to artefacts) and, thus, predicted by Bird et al. (2000) to be selectively impaired along with artefacts hitherto have not been reported in the literature.

Motoric features. Again, fine-grained differences between semantic fields were observed. Within words referring to objects, vehicles (27.6% of all features) were the most dependent upon motoric features, body parts (17.4%) and miscellaneous artefacts (15.4%) moderately so, along with fruits and vegetables (13.5%), where motoric features were related to food preparation e.g. <peel>, <cut>. Other fields were less dependent on motoric features (clothing for example, had virtually no motoric features: 1.5%). Vehicles had significantly more motoric features than body parts ($t(40) = 3.01$, $p = .002$), body parts had more than fruit and vegetables ($t(52) = 2.380$, $p = .010$), fruits and vegetables had more than animals ($t(57) = 2.980$, $p = .002$) and animals had more than clothing ($t(43) = 4.652$, $p < .001$). In contrast to objects, words referring to actions were far more dependent upon motoric features ($t(454) = 15.182$, $p < .001$). Among the action fields, especially, manner of motion (56.4%), but also change of state (41.1%), body action (36.3%) and tool action (38.0%) were highly dependent on motoric features. Fields of communication (15.0%) and cooking (16.5%) were moderately dependent on motoric features, and a few action fields (e.g. light emission at 6.6%) had hardly any motoric features.

Overall, the present findings are consistent with the claims of accounts that assume differences across concepts in the number and weight of different feature types (e.g. Cree & McRae, 2003; Farah & McClelland, 1991; Tyler et al., 2000). The investigation in the present study, however, went a step further than previous studies by making distinctions on a more fine-grained level by distinguishing between Visual and Other Perceptual features, and also between Functional and Motoric features. These distinctions are important with respect to predicting fine-grained patterns of performance. Although selective impairments for living things relative to artefacts have been claimed to be associated with the combined effects of impairments to visual *and* other sensory features, different patterns of performance may result depending upon which sensory classification is used. For example, animals and fruit/vegetables are indistinguishable when visual and other sensory features are combined, but are distinguished when only visual features are considered, animals being more dependent on visual features than fruit/vegetables. The fine-grained distinction here allows dissociations between animals on one hand and fruit/vegetables on the other, which could explain cases in which one of these fields is spared but not the other (e.g., Hart et al., 1985).

The present data also provided novel predictions with respect to what type of impairments we should expect in the action domain, depending upon the type of features that make up the semantic organisation of the different words referring to actions. While motoric features were shown to be important for most of the semantic fields in the domain of words referring to actions, some action fields could be distinguished on the basis of feature types. For example, words referring to light emission were highly dependent on visual features, while noise making and cooking words were dependent on sensory features from other modalities. Interestingly, these distinctions in the domain of actions show a degree of similarity to the sensory and functional distinctions in the domain of objects. Therefore, assuming the sensory-functional theory of naming impairments, "category-specific" effects should also be observed for words referring to actions, provided that suitable items are used for testing. It remains to be seen whether the absence of reports of such cases in the literature is merely a consequence of the lack of attention of past research to impairments in the domain of actions. Alternatively the lack of selective deficits within the action domain may constitute evidence against the sensory-functional theory (for additional discussion and simulation results consistent with the analyses above, see Vinson & Vigliocco, 2002; Vinson et al., 2003).

Relationships between words

The previous analyses were concerned with the feature properties of individual words. Here, relations between words, as illuminated by their featural makeup, are explored, focusing upon a number of dimensions about which specific claims have been made. One set of analyses consider shared features and correlated features (e.g. Devlin, Gonnerman, Andersen & Seidenberg, 1998; Garrard et al., 2001; McRae et al., 1997; Tyler, Moss, Durrant-Peatfield & Levy, 2000), and ask whether words referring to objects and words referring to actions, and living and non-living things, differ along these dimensions. A second set of analyses consider the distinctiveness and correlation of features in the fields of animals and tools, testing specific claims made by Tyler and colleagues (Tyler et al., 2000) about the types of features that tend to be intercorrelated.

Shared features. Traditional featural views of semantic representations (e.g. Norman & Rumelhart, 1975; Rosch & Mervis, 1975; Smith, Shoben & Rips, 1974) and current work (e.g. Maki, Krimsky, & Muñoz, 2006; see Lucas, 2000; Hutchison, 2003 for reviews of the semantic priming literature) use some kind of feature overlap as the crucial measure of similarity between words' meanings. At least for some semantic fields of nouns referring to objects, shared features have been shown to have an effect on performance in behavioural tasks such as semantic priming (McRae & Boisvert, 1998). The notion of shared features is also important to many views of semantic representation and impairment, even if they do not specifically rely upon shared features *per se* as the centrally important aspect of meaning. For example, theories such as the Conceptual Structure account (Tyler et al., 2000), and other theories in which intercorrelation or distinctiveness among features is important (e.g. Devlin et al., 1998; McRae et al., 1997), depend upon shared features as this is the only way intercorrelation can arise. Moreover, feature distinctiveness is defined in contrast to shared features: distinctive features are those that are not shared among many exemplars.

The present study thus begins with analyses of shared features. The raw similarity among words within the complete featural space was assessed by considering the extent to which words had features in common. First, a measure of shared features was calculated on the basis of the number of features shared between two words (see the first column of Table 2).

Table 2. Average number of shared features and shared feature weights (standard deviations in brackets) as a function of semantic field classification for words referring to objects and actions (for actions, semantic field labels are taken from Levin, 1993).

Field	Mean shared number of features	Mean shared feature weights
Animals	7.78 (3.36)	27.73 (11.20)
Fruit & vegetables	9.38 (2.91)	34.88 (15.11)
Tools	8.80 (3.39)	31.66 (13.58)
Vehicles	9.44 (4.10)	27.14 (13.04)
Body parts	6.56 (4.05)	18.84 (14.17)
Clothing	8.48 (3.43)	31.39 (12.20)
Misc. artefacts	5.65 (3.61)	12.17 (9.91)
ALL OBJECTS	**3.00 (2.50)**	**7.26 (8.48)**
Body actions	4.47 (2.24)	11.34 (5.86)
Body sense	5.52 (2.84)	14.77 (8.55)
Change of location	6.44 (2.92)	17.91 (7.53)
Change of state	5.60 (2.60)	13.84 (6.59)
Noises	6.52 (3.09)	20.07 (9.47)
Communication	6.70 (2.79)	16.94 (7.67)
Construction	8.71 (3.65)	23.24 (9.91)
Contact	8.60 (4.09)	26.94 (14.16)
Cooking	6.90 (2.92)	28.86 (13.21)
Destruction	8.57 (4.11)	20.00 (9.05)
Exchange	7.07 (2.76)	20.63 (8.67)
Heat/light emission	4.65 (2.80)	15.71 (11.46)
Motion direction	4.53 (1.80)	12.27 (4.89)
Motion manner	5.75 (2.20)	16.75 (6.70)
Tool action	8.44 (3.58)	18.75 (8.31)
ALL ACTIONS	**3.33 (1.64)**	**8.16 (4.24)**

Substantial variability was observed. For words referring to objects, the number of features shared among exemplars in a semantic field was numerically largest for fruits/vegetables, tools and vehicles and smallest for animals, body parts and miscellaneous artefacts (although it should be noted that these differences do not reach statistical significance, a consequence of wide variability among items, stemming from the varying properties of similarity within a given category). These tendencies may reflect the existence of subfield distinctions in these latter fields such as, for example, wild vs. domestic animals, face vs. limb-related body parts, and furniture vs. buildings. In the action domain, the largest number of features were shared in fields of contact, construction, destruction and tool action, and the smallest number was shared among body action, heat/light emission and direction of motion. Importantly, comparing the domains of objects and actions, there were, roughly, similar numbers of semantic fields with many and few shared features (although fields with the fewest shared features were in the action domain), although considering objects and actions without regard to semantic field distinctions it is interesting to note that these objects tended to share less features than actions ($t(383) = 6.75$, $p < .001$), perhaps illustrating a general tendency for object-nouns to be organised into separable categories while action-verbs tend to share features more generally.

In order to assess featural overlap taking weights into account, it was first necessary to create an index of shared weights between word pairs, which was calculated as follows. The weighted feature overlap measure for a given pair of words was defined as the sum, across all features, of the minimum feature weight for the two words, taking into account only those features with nonzero weights. Since only the total value of featural weight is taken into account (and not the number of features contributing to this measure), the measure is the same for two words sharing ten features with weights = 1, and for two words sharing only one feature with weight =

10). Shared (weighted) feature overlap was first assessed within semantic fields, to illustrate the characteristics of semantic fields as above (see the second column of Table 2 above).

Considering words referring to objects, the highest feature weight similarity was again found among fruit/vegetables, tools and clothing, with lowest similarity levels among miscellaneous artefacts and body parts, although again these differences do not reach statistical significance due to the extremely high variance at the item level. Considering words referring to actions, the highest weight similarity was observed among words referring to cooking, contact and construction, with lowest similarity among direction of motion, body action and change of state. Words referring to objects had higher within-field shared weights than words referring to actions ($t(383) = 12.51$, $p < .001$), although this may be a consequence of the higher weights for objects overall as discussed earlier in this chapter. Further, this pattern is reversed when semantic fields are disregarded and shared weights are considered within all objects and within all actions ($t(383) = 6.75$, $p < .001$): feature weights tend to be shared more widely across action-verbs while shared weights tend to be limited to members of the same superordinate category in the object domain.

These patterns of similarity do not differ drastically whether feature weights are considered or not, at least when comparing semantic fields to each other, suggesting that they are quite robust in indicating a hierarchy of differences in shared features within semantic fields, possibly, illustrative of the relative semantic density of the fields investigated.

<u>Correlation among features</u>. The pattern of correlation among features is closely related to featural overlap, because features that frequently overlap will be highly correlated to each other. Correlation, however, takes not only overlap into account, but also the relative amount of overlap for pairs of features vs. those instances in which

only one of a pair of features occurs for a given item. Analysis of correlation first considers how strongly correlated features are for different semantic fields, but can also be applied to questions related to specific types of features (whether different broad classes of features tend to co-occur depending upon the semantic field in question). The Conceptual Structure account by Tyler and colleagues (2000) makes specific claims about correlation, suggesting that living things have more strongly intercorrelated features than non-living things, and words referring to objects have more strongly intercorrelated features than words referring to actions. This is consistent with the analyses of shared features reported above: living things tend to share more features than artefacts, which in turn tend to share more features than action-words. Tyler et al. also make specific claims about the kind of features that are correlated, depending upon semantic domain. Sensory features of living things are, in general, less distinctive, and thus the strongest correlations between features are expected to be between sensory features shared among many exemplars, such as the fact that things that have tails also tend to have legs. Such features are typically strongly correlated (often with numerous other features as well) and shared among a number of exemplars). For artefacts, on the other hand, sensory features tend to be distinctive, and thus should be correlated not with other sensory features but with their related functions, such as the fact that things that are sharp are also used for the function of cutting, a correlation much stronger than between sensory feature <sharp> and other related sensory features like <hard>. However, counterevidence has been provided by Garrard et al. (2001) based on analyses of speaker-generated features. According to Garrard et al., the greatest degree of intercorrelation among features is observed for distinctive features of living things (e.g. beaks and wings are quite distinctive when considering all living things, but there are virtually no exceptions to their co-occurrence), higher than either shared or distinctive features for artefacts. Garrard et al. further showed that the representations for artefacts are not strongly dependent upon distinctive form-function correlations.

It is premature, however, to discount the Conceptual Structure account based on this evidence, because some criticisms may be levelled against the feature norms obtained by Garrard et al. (2001). First, they were obtained using a rather restrictive methodology in which participants were asked to generate a certain number of features of a given kind for each word in the set by completing phrase frames (six features for each of the phrase frames "IS ____", "HAS ____" and "CAN _____", plus one "category" feature, in which participants were meant to report a superordinate category). This method of feature collection, which was identical for all items, may have led participants to generate predominantly shared features for category exemplars. Second, Garrard et al. used a limited set of words, of eight semantic fields, and eight exemplars in each field. Third, all their items were labels of concrete objects. This point is important as the specific context in which feature generation is employed (i.e., the other words included in a feature generation list) could, in principle, bias participants to produce only those features sufficient to distinguish a given item from others in the set, which could be a crucial weakness of this methodology (see Murphy, 2002). It is therefore important to replicate Garrard et al.'s tests of the predictions of the Conceptual Structure account, using different methods of feature collection less susceptible to these criticisms. The methodology in the present study differed from those of Garrard et al. (2001) as a wide range of semantically unrelated exemplars (including numerous words referring to actions) were included and the participants were allowed to generate whatever type of features they felt were important to the describe the meaning of a given word rather than being constrained to produce only certain kinds of features, and in certain proportions.

In order to provide a close basis for comparison, these analyses focused upon a limited set of exemplars from three semantic fields: animals, tools and action words. A limited set of words was investigated in this case to allow equating the words for concept familiarity, an important consideration because less familiar items may exhibit less typical featural profiles (for example, participants might only produce generic features related to a superordinate if the item itself is not so familiar). Items used for this comparison were 12 animals (*bird, camel, cat, dog, fish, fox, goat, horse, lion, mouse, sheep, tiger*), 12 tools (*fork, tweezers, brush, pencil, pen, pliers, chisel, scissors, razor, gun, file, hammer*) and 12 verbs referring to actions (*touch, clang, smell, hold, throw, frown, drill, write, twist, spray, exchange, inhale*).

Following the analysis used by Garrard et al. (2001), first the value of the correlation coefficient for all possible pairs of features across exemplars (in the entire set of 456 words) were calculated (taking feature weight into account). For each of the words in the test set (12 animals, 12 tools, 12 actions), all possible pairings of that word's features were assigned average correlation values based on the correlations in the entire set of 456 words (e.g. if <wings> and <beak> had a correlation of $r = +0.96$ across all items in the set, this value was assigned to the feature pair <wings>, <beak> for the item *bird*.). These values excluded the numerous instances where neither feature occurred for a given word, which would have produced inflated correlation measures due to the sparseness of the feature vectors. These values were then averaged across all feature pairs for a given word. The average correlation coefficient for each semantic field was 0.146 for animal features, 0.119 for tools, and 0.081 for actions; pairwise comparisons between feature pairs (items as a random factor) using nonparametric tests (Mann-Whitney U) revealed that all three correlations differed significantly from each other (all $p<.001$). This is consistent with other analyses of featural correlations within object domains (e.g. Garrard et al., 2001; McRae et al., 1997) in which the features of animals are more highly intercorrelated than the features of tools, for which distinctiveness is argued to be more important. These results also provide novel information about the action domain for which correlations are overall lower than for objects despite the overall tendency for actions to share features more than objects.

Distinctiveness and correlation. Differences between the semantic fields of animals and tools in terms of feature distinctiveness were assessed following Garrard et al. (2001). A feature's distinctiveness was operationalised as being the proportion of words within a semantic field that share that feature (weights > 0). Thus a value of 1.0 indicates a feature that is shared among all exemplars within a field, and smaller values indicate higher level of distinctiveness (features with values of zero were excluded, as they are by definition not representative of any exemplars in a field). Features with values greater than 0.5 on this scale were considered to be "shared" and those 0.5 or below, "distinctive". To assess whether living things were predominantly characterised by shared and correlated sensory features, and artefacts by distinctive features (for which form and function are correlated, such as <sharp> and <cut>), the following analysis was performed. The dependent measure was the number of statistically significant correlation coefficients between pairs of features for each word. Two different types of feature pairings were considered: "intracorrelation" (as described by Garrard et al.), the pairing between two features of the same type (sensory-sensory or functional-functional); and "intercorrelation", the pairing between features of different types (sensory-functional). Feature pairs were also divided into two conditions according to their overall distinctiveness (shared vs. distinctive features), thus allowing a 2x2 factorial analysis of pairing type and distinctiveness. Separate analyses were carried out for animals and tools. The overall proportion of statistically significant correlations (correlations that differed from zero, alpha = .05) in this set was very low (animals =

9.1%, tools = 8.4%, actions = 8.0%); proportion of significant correlations by condition is reported in Figure 3.

Figure 3. Average proportion of feature correlations that were statistically significant for animals, tools and actions as a function of featural distinctiveness and feature-correlation type, considering only sensory and functional features. Error bars indicate standard error of the mean (by feature pairs).

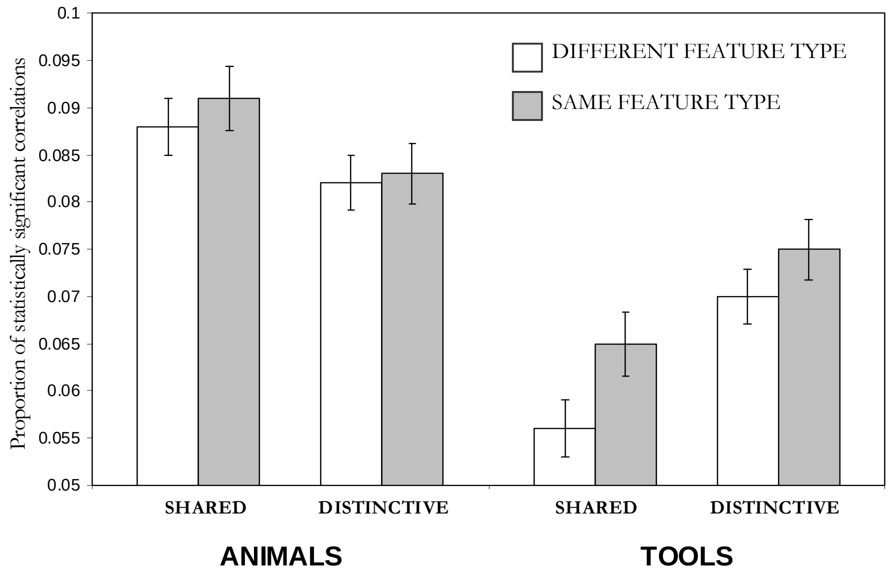

Animals had a small but significant tendency to have more correlation involving shared features than distinctive features ($F(1,11) = 5.21$, $p = .03$); however the effect of correlation-type was not significant, nor did the two factors interact (Fs < 1). For tools, instead, main effects of distinctiveness (more distinctive features than shared; $F(1,11) = 7.36$, $p = .020$), as well as a main effect of correlation-type (more correlations within features of the same type than across feature types; $F(1,11) = 6.90$, $p = .024$), but no interaction between the two ($F < 1$), were observed. This replicates the general pattern found by Garrard et al.--more correlation among shared than distinctive features for animals, more correlation among distinctive than shared features for tools--but runs counter to the prediction that form-function correlations (intercorrelations) should be more prevalent in artefact domains. In addition to this contrast, animals and tools were directly compared, by comparing distinctive features for animals to distinctive features for tools. Distinctive features for animals were more likely to be significantly correlated (8.3% of cases) than those for tools (7.3%), a significant difference ($t(22) = 2.49$, $p = .021$).

To summarise, in the analyses reported above, some--but not all--of the differences that have been claimed to be important in determining concept organisation for different fields were observed. Within the object domain, living things differed from non-living things with respect to correlated features (more common for living than non-living) but not with respect to the number of features and the number of shared features, as has been argued by other accounts (e.g. Tyler et al., 2000). Furthermore, despite methodological difference in feature collection methods, these analyses replicated the general finding by Garrard et al. (2001) that distinctive features of living things were more correlated than features (distinctive or not) of artefacts, contrary to the predictions of Tyler et al. (2000). Hence, this work does not support Tyler's account of category specificity.

Finally, contrasting the domains of objects and actions, as suggested by Huttenlocher and Lui (1979), the two were found to differ along the dimensions of feature numbers (richer representations for words referring to objects than words referring to actions) and proportion of correlated features (more for objects than for actions). However, objects and actions differed with respect to number or weight of shared features, depending on whether the analysis was fine-grained (considering shared features within semantic fields, where objects exhibited more shared features than actions) or more coarse-grained (ignoring semantic field differences, objects exhibited less shared features than actions overall). These findings emphasised the importance of differences in terms of correlated features between domains, and highlight the possible problems in considering only a single measure of similarity such as shared features, which may not adequately characterise the true semantic relations among items in a set.

Chapter 4: Modelling the lexical-semantic level of representation

The featural representations described in the previous chapter can perhaps best be described as representing conceptual knowledge - nonlinguistic mental representations of things, events, etc. However, as discussed in Chapter 1 there are compelling reasons to posit a distinction between conceptual and semantic representations. A way to conceptualise this distinction is to assume that only concepts have featural representations, and the lexical-semantic level of representation binds featural representations to serve language functions. This level, at least for concepts that are lexicalised in a language, serves to mediate between meaning, syntax and wordform. Representational architectures of this kind have been described in neural terms by Damasio et al. (2004) as convergence zones, which connect different brain areas responding to different streams of sensorimotor information. Such an architecture applied to the meanings of words would avoid the problems associated with theories which do not include a conceptual/semantic distinction (or which do not include separate organisation at these levels): only those concepts which are lexicalised would have representations at this lexical-semantic level, which could differ cross-linguistically. Further, the language-specific effects observed only in verbal tasks (e.g. Brysbaert et al., 1998; Vigliocco et al., 2005; Kousta et al., in press) would arise at this lexical-semantic level of representation, leaving conceptual representations unaffected by linguistic differences.

Conceptual basis of modelling

In the present study, properties of the speaker-generated features provide the basis for modelling semantic organisation. Note that this is importantly distinct from the features themselves, which are argued to represent conceptual information; properties of features concern their co-occurrence, patterns of correlation, and other aspects concerning their combination into the meanings of words. This general framework will henceforth be termed **FUSS**: Featural and Unitary Semantic Spaces: a conceptual representational space which is operationalised by the speaker-generated features themselves, and a semantic representational space derived from properties of similarity among words' featural content.[8] This latter space is obtained using a technique which does not require making *a priori* assumptions about which specific properties of features are responsible for characteristics of organisation at this level: self-organising maps (SOMs, Kohonen, 1997). This approach takes a multidimensional input and produce as output a lower-dimensionality space in which important relationships among entities in the input are preserved.

Representing meaning with self-organising maps

Self-organising maps are particularly well-suited for modelling lexical-semantic representations because they create a spatially-organised output network ("map") of units in an unsupervised manner based upon various differences between different representations in the input space without the need to identify which particular aspects of the input are important in determining similarity (Kohonen, 1997). Important elements of self-organising maps are first, the input vectors. Each input vector corresponds to a single concept to be represented by the output network, made up of numeric entries in an input space of high dimensionality. Sets of input vectors make up the training set, which represent the experience by which the model learns its representations. The output map is a low-dimensionality similarity space, intended to ultimately reflect similarity structure in the input space. Crucially, each unit in the output network is associated with two types of information: co-ordinates of its spatial location in the output map, and a prototype vector of equal dimensionality to the input. This

[8] Reference in the text to the Unitary Semantic level of representation will often refer to "lexical-semantics" as a reminder of the assumption that this level is separate from conceptual representation.

serves the dual roles of vector projection (nonlinearly projecting an input space of high dimensionality to an output space of much lower dimensionality) and vector quantisation (providing a point estimate for a region in the input space). This is achieved via a training regimen which allows the system to develop through experience, adjusting the spatial coordinates on the output map so that concepts with similar input vectors appear on nearby regions on the map. Each training event consists of presentation of an input vector, which is compared to all prototype vectors. The one unit whose prototype vector is most similar to that input (according to Euclidean distance) is selected as the "winner", and the values of that prototype vector are adjusted in the direction of the input vector according to some function (typically a decreasing function which applies the greatest amount of change at the early stages of training and reduces to a small amount over the course of training). These adjustments, however, are not limited to the winning unit, but also extend to its spatial neighbours on the output map, the extent of adjustment depending upon a neighbourhood function which, again, typically reduces in size over the course of training. This reducing neighbourhood function serves to organise the output space coarsely in the early stages of training (reflecting the greatest regular differences in the input), and then gradually narrows its focus, thus reflecting finer degrees of similarity as training proceeds. Finally, each self-organising map is characterised by an initial state; the initial prototype vector for each location in the output map is initialised to a starting value. This reflects the state of affairs before any training has occurred. Typically this is done either by assigning random weights, or by applying some regular function related to two-dimensional position.

In this case, the multidimensional input is the 1029-dimension space defined by the number of features, with one vector for each of the 456 words in the feature set (thus each word can be considered as a point in a 1029-dimensional conceptual representation space), A two-dimensional output layer is taken to reflect the lexical-semantic level of representation. Application of the SOM algorithms results in local clustering of words with shared features, emphasising those properties of features that are most crucial for distinguishing between features. Organisation occurs not only on the basis of whether words share features or not, but also on other properties of features such as the weights of features (higher weighted features have greater impact than lower weighted features), the distinctiveness of features (features shared by a large number of words have less impact than features shared by fewer words, although widely-shared features will have greater global impact), and co-occurrence of the features (features correlated with each other will offer mutual support to the development of the output map, while features that never co-occur may have opposite effects, not only upon each other, but also upon other features correlated to one but not the other). In short, the trained SOM output space reflects the combined influence of a number of properties of featural representations, and these properties need not be specified in advance (or even known).

One concern about the use of SOMs involves the reduction of a space of high dimensionality into one of lower dimensionality is that the lower-dimensionality representation may result in coincidental proximity among concepts, especially because equal distances on the spatial map may not correspond to equal distances in prototype space. For example, consider the simple example of a linear sequence (1-2-3-4-5-…-n) mapped into two dimensions. Such a sequence can successfully be mapped into the bi-dimensional space in any number of ways, provided that adjacent values are still adjacent in the final map. Figure 4 (see next page) illustrates two such cases.

Figure 4. Two possible configurations of the numeric sequence {1 ... 12}
as it could be represented in a two-dimensional output space (4 x 3)

1	2	3	4
8	7	6	5
9	10	11	12

12	11	2	3
9	10	1	4
8	7	6	5

In both cases, the sequence was exactly preserved, but the constriction of the sequence into a two-dimensional space produced misleading conclusions about similarity of concepts in the space. In the example on the left, one might conclude from the spatial arrangement alone that 1 is equally similar to 2 and 8, more similar to 7 than to 3, etc. In the example on the right, instead, 1 is not at all similar to 8 but is instead equally similar to all members of the set {2, 4, 6, 10}. Such concerns can be reduced by starting with a large neighbourhood radius and gradually reducing the size of the training effects, thus serving to initially provide a very coarse-grained arrangement which reflects the most distinctive differences among the items in the training set (in the cases above, largest numbers vs. smallest numbers), and eventually capturing the finer qualities of the input set (e.g., local proximity between values adjacent on the number line). But this does not entirely eliminate such problems, especially when we consider cases more complex than the ordered numeric set illustrated above. However, if multiple SOMs are trained using the same input vectors, but using different (random) starting configurations of the output space, such coincidental (misleading) proximities will be reduced, while, instead, truly proximal concepts will remain proximal across output maps. Considering the examples in Figure 4 above, merely averaging across the two maps is enough to remove most of the coincidental proximities (e.g. the only immediate neighbours in common for 2 between the two maps are 1 and 3; the only common neighbours of 8 are 7 and 9) although a few instances still remain (e.g., these two maps together are not enough to rule out 7 and 10 as immediate neighbours). In order to avoid problems of this nature in estimating semantic similarity in feature-based maps, ensemble average distances were calculated: the similarity of two words was determined by the average of their Euclidean distances across multiple output maps. This allows a means of preserving the regular relations based upon featural contrasts and similarities (see Kohonen, 1997), and is analogous to averaging across speakers, each of whom has idiosyncratic relations among lexical concepts based on personal experience, but who share overall commonalities based on common reference and linguistic convention.

Method

Input data consisted of the 456 x 1029 (word x feature) weight matrix, and each of 100 output maps was defined as a rectangular space of 40x25 units, arranged in a rectangular lattice. This dimensionality was determined based on the two principal eigenvectors of the input data vectors, and number of units based upon the number of input vectors (Kohonen, 1997). Each unit on the output map is associated with a 1029-dimension prototype vector associated with a (two-dimensional) location on the output map. These prototype vectors were initialised to random values, independently for each output map. Training was conducted using SOM Toolbox 2.0 (http://www.cis.hut.fi/somtoolbox/) which implements self-organising maps (Kohonen et al., 1996) within MATLAB. There were two steps in the training of each map: first a "rough" step intended to be sensitive to the most salient distinctions, followed by a "fine" step to make more sensitive adjustments at a local level. The rule for adjusting weights is as follows:

$$\mathbf{w}(\mathbf{t}+1) = \mathbf{w}(\mathbf{t}) + \mathbf{a}(\mathbf{t})\mathbf{K}(\mathbf{t})\ [\mathbf{x} - \mathbf{w}(\mathbf{t})]$$

where $\mathbf{w}(\mathbf{t})$ is the weight at time $\mathbf{t}$, $[\mathbf{x}(\mathbf{t}) - \mathbf{w}(\mathbf{t})]$ is the difference between vector $\mathbf{w}$ and the input vector $\mathbf{x}$ (constant over $\mathbf{t}$), $\mathbf{a}(\mathbf{t})$ is the learning rate (here, defined as a linear decreasing function, in the rough phase beginning at 0.5 and in the fine phase beginning at 0.05, with an intercept of zero at $\mathbf{t}_{total}+1$ where $\mathbf{t}_{total}$ is the number of training epochs in a given phase: 40 for the rough phase and 160 for the fine phase), **K**(**t**) is the Gaussian neighbourhood kernel function

$$\mathbf{K = \ exp\ (-d^2 / 2\sigma(t)^2\ )}$$

where **d** is the distance between a unit and the winner, and **σ**(**t**) is the neighbourhood radius at point **t** in time. Neighbourhood radius **σ** also decreases linearly: in the rough phase from 20 (half of the maximum map dimension) to 10, and in the fine phase from 10 to 1.

Each map was trained in batch mode, which means that all feature vectors were presented to it in a single epoch (rather than presenting one of the feature vectors and then adjusting the map based only on that item), and all adjustments based on the identification of the "winner" for each input vector and the application of the neighbourhood function were simultaneously applied at the end of a set of epochs (comprising one presentation of each feature vector). In essence, this means that each prototype vector is replaced with its weighted average over each of the input samples (i.e. one vector corresponding to each of the 456 words in the training set), where weights are assigned according to the neighbourhood function **K** above (given that the learning rate **a**(**t**) is constant for a single epoch). The use of batch mode thus minimises idiosyncratic fluctuations of the output map based on presentation order of the individual input vectors, and thus requires fewer training cycles to reach a stable configuration, similar to the end configuration resulting from randomly-ordered presentation of single vectors. It also has a significant advantage in the amount of time/computing resources required; batch mode requires only one application of the neighbourhood function for the entire set of words, vs. 456 applications if each word is presented individually.

Once the fine training phase was complete, the output map was considered to be complete.[9] Each map was then labelled with the words in the training set: whichever unit was most similar to a given input vector was labelled with the corresponding word. Euclidean distances between all possible pairs of words were calculated, giving a measure of word-word similarity for that map[10]. Once this was complete, a composite distance measure was obtained, by averaging distances for each word pair across all 100 maps. This distance measure serves as an estimate of semantic (dis)similarity based on characteristics of the featural input, and thus provides the basis for evaluating the lexical-semantic representations arising under the FUSS model as described above. Global and local properties of this lexical-semantic space will be discussed in the next chapter.

[9] This procedure was judged to be sufficient based on visual inspection of the first few maps created with these parameters, which appeared to suitably reflect semantic similarity, at least according to intuition. In subsequent chapters this will be assessed more formally.

[10] In a few cases, a single unit in the output map could be the best matching unit for more than one input vector. In this event, that unit was given multiple labels, and the distance between them was zero. Such cases can be considered "true synonyms", i.e. indistinguishable according to that particular map.

Chapter 5: Properties of the lexical-semantic space in FUSS.

Given the transformation of the speaker-generated features into a composite semantic similarity space as described in the previous chapter, it is first necessary to establish that the resulting representations actually do reflect semantic similarity. In other words, assessing whether words with similar meanings are close to each other, and words with dissimilar meanings are far apart. This serves as the most basic test of FUSS, because any acceptable model of semantics must be able to capture the gross distinction between related or unrelated words. Since independently-obtained measures of semantic similarity for the items included in the present set were unavailable, it was necessary to start by assessing whether the organisation of this space conforms to intuition about which items should be similar to each other and which should not (this will be complemented with behavioural evidence in subsequent chapters). This question about organisation will be examined at a number of levels of specificity, ranging from similarity among words, to similarity among semantic fields, to similarity across the object and action domains. Analysis of the clustering performance of nouns referring to actions will also reveal the relationship between grammatical class and semantic representation. This is an important question because many studies investigating grammatical class distinctions between nouns and verbs have conflated this grammatical class distinction with the semantic distinction between objects and actions (see Vinson & Vigliocco, 2002; Vigliocco, Vinson, Arciuli & Barber, 2008).

Contrasting semantic fields

The first analyses tested whether sets of words of the same (intuitively designated) semantic field are near each other and distant from members of other semantic fields in FUSS lexical-semantic representation space. These analyses also indirectly provide information concerning proximity among semantic fields (i.e., groups of words related in meaning according to the similarity space) to evaluate whether the clustering patterns between semantic fields reflect (intuitive) semantic similarity. These analyses were conducted separately for words referring to objects and verbs referring to actions (for now, excluding nouns referring to actions).

Here, the average semantic distance between words from a given field (e.g., all possible pairings involving two words referring animals) was compared to the average semantic distance between those exemplars and exemplars from other fields (e.g., all possible pairings involving one word referring to an animal, and one word from another semantic field). If FUSS semantic similarity measures reflect this kind of category-level similarity, the within-field distances should be much less than cross-field distances. Distances for object fields are reported in Table 3.

Table 3. OBJECTS. Average distance between exemplars of the same semantic field, and between exemplars of different semantic fields (standard deviations in brackets). Distances are measured in arbitrary units based on ensemble averages of maps with dimensions 40x25 units. All within vs. between comparisons are significant using independent-samples t-tests; $p < .001$(one-tailed).

	Within-field distance	Between-field distance
Semantic field:		
fruit/vegetable	7.0 (3.4)	23.0 (1.7)
body parts	16.0 (6.0)	18.7 (2.7)
animals	6.1 (2.9)	20.5 (1.6)
clothing	4.2 (1.8)	20.7 (1.2)
tools	9.9 (3.4)	19.5 (2.2)
vehicles	8.6 (5.8)	21.0 (2.3)
other artefacts	12.0 (4.8)	19.3 (3.1)

In all of these cases, the average distance between exemplars of a single semantic field was significantly lower than distances between exemplars of different semantic fields, suggesting that distances are capturing properties of family resemblance associated with these semantic fields. This was so even for less-coherent fields. For example, body parts were highly segregated into subfields, distinguishing between facial parts and limbs/extremities, which had very few features in common. Similarly, "other artefacts" included items of furniture; parts of buildings like *wall, ceiling, floor*; and other items such as *bomb, book, box* which are difficult to classify at a finer level. Nonetheless, these items were overall more similar to each other than to exemplars from other object semantic fields. The same analysis was carried out for action fields, as reported in Table 4.

Table 4. ACTIONS. Average distance between exemplars of the same semantic field, and between exemplars of different semantic fields (standard deviations in brackets). Distances are measured in arbitrary units based on ensemble averages of maps with dimensions 40x25 units. All within vs. between comparisons are significant using independent-samples t-tests; $p < .001$(one-tailed).

Semantic field:	Within-field distance	Between-field distance
body action	16.4 (4.1)	18.1 (3.6)
body sense	14.6 (3.6)	18.4 (2.4)
change location	7.7 (3.2)	15.4 (1.9)
change state	8.4 (3.7)	15.0 (2.1)
communicate	9.8 (3.5)	16.7 (2.8)
construct	5.9 (2.3)	16.0 (1.4)
contact	6.5 (3.4)	16.3 (1.2)
cook	5.4 (3.2)	19.5 (1.0)
destroy	9.3 (4.4)	15.4 (2.4)
exchange	5.7 (3.3)	16.7 (1.5)
light emission	9.2 (6.1)	19.1 (1.6)
motion direction	7.9 (2.6)	15.3 (1.8)
motion manner	11.0 (3.7)	17.0 (2.3)
noise	6.1 (2.5)	19.6 (1.5)
noise animal	4.0 (2.1)	20.7 (0.7)
tool action	10.4 (4.0)	17.2 (2.2)

As was the case for the object fields, within-field distance measures were significantly different from the between-field measures for every field listed. Again, this was even the case for the less coherent fields such as 'body actions', a somewhat generic semantic field encompassing verbs such as *inhale, inject, itch, lick, retch, sit, spit, wash*; or 'body senses', including disparate sensory verbs such as *feel, hear, listen, look, smell* (see Appendix A for a full set of items and their semantic field labels), reflecting the fact that even these kinds of words were somewhat semantically clustered, and roughly separable from other action verbs as a whole.

Interestingly, the between-field distances for words referring to actions were notably smaller than between-field distances for words referring to objects. This is because action fields tended to be close to each other, while this was hardly ever the case for object fields (also illustrated in the analyses of shared features and shared feature weights reported in Chapter 3). For example, noises and animal noises were separated by an average of only six units, change of location and direction of motion by only nine, change of location and change of state by 10, tool action and construction by 10, manner and direction of motion by 11, destroy and change of state by 12, exchange and change of location by 12, animal noise and communication by 15. For object fields, the most similar classes were body parts and clothing (15), tools and miscellaneous artefacts (16), animals and body parts (17), vehicles and tools (18). These relationships among relatively proximal semantic fields make intuitive sense, but this will be formally tested (at least for certain semantic fields) in the behavioural experiment reported in Chapter 8.

Figure 5 (see next page) illustrates the general tendency for object semantic fields to be more strongly differentiated from each other than action semantic fields. A clear category division is observed between "fruits and vegetables", while actions involving exchange blend into "communication" (exchange of information) which in turn blend into "manner of communication", which in turn blend into "non-communicative sounds". Importantly, in both cases local proximity remains consistent with intuitive judgements of similarity.

Figure 5. Two-dimensional projection of semantic similarity space for fruit and vegetables (left panel) and selected words referring to actions (right panel). Distances between words reflect degree of semantic (dis)similarity.

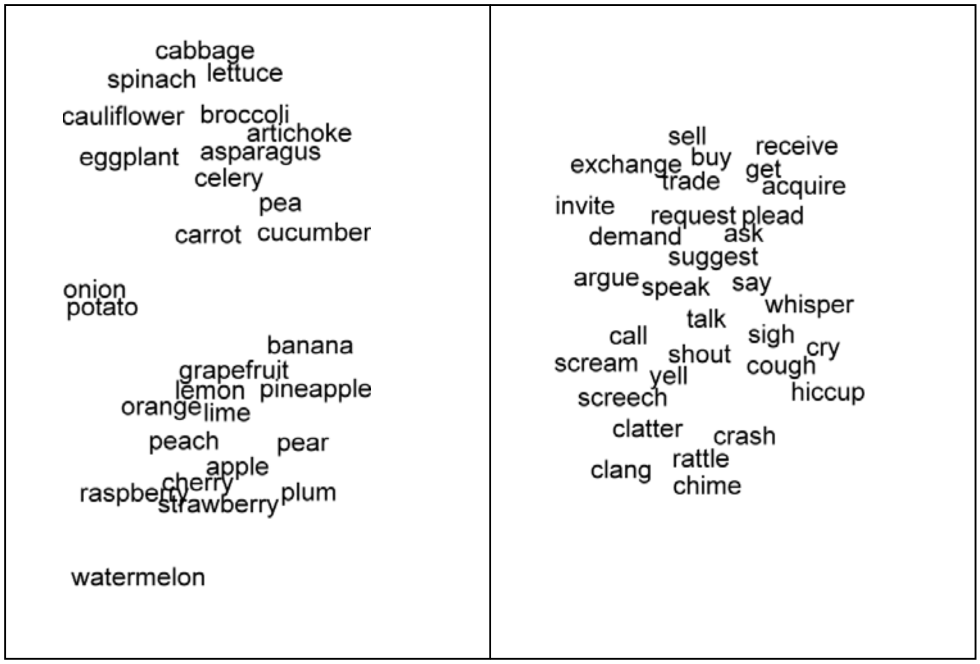

The different patterns of similarity for words referring to objects and words referring to actions are consistent with claims that have been made about the differences between the lexical-semantic organisation of objects and actions as discussed in Chapter 1. The words referring to objects in the model are organised categorically, with few intermediate exemplars (for example, onions and potatoes, which seem to form a separate subcategory from the majority of other vegetables, and watermelon which is somewhat distinct from other fruits), while category boundaries were shown to be essentially meaningless for the words referring to actions depicted above.

Contrasting objects and actions

Another crucial question is the extent to which words referring to actions and words referring to objects are separable in FUSS lexical-semantic similarity space. Because objects and actions are different in so many ways (e.g., objects refer to an identifiable entity and actions express inter-entity activities; they fulfil different sentential functions; they differ in interconceptual organisation, Graesser et al., 1987; Huttenlocher & Lui, 1979) it is important to demonstrate that they are also separable in lexical-semantic space. Were this most basic distinction not reflected in the organisation at this level, it would cast doubt upon FUSS even if it proves able to capture fine-grained similarity.

In a first analysis, semantic field proximities were calculated across the object/action divide, with special attention paid to those object and action fields that were shown to be close to each other. Fruit/vegetables were close to cooking (average distance 15 units), and tools were close to tool actions (11 units), construction (12 units), and destruction (16 units). These proximities were almost entirely due to item-specific proximity between words referring to tools and their most closely associated actions (e.g. *to shovel* was close to *the hoe* and *the shovel*; *to drill* was close to *the drill*, etc.), but almost none of these actions was as close to its most associated object than that object was to other related objects (e.g. *the hoe* was 4.4 units from *to hoe*, but only 2.4 units from *the broom* and *the rake*; *the drill* was 3.1 units from *to drill*, but only 2.2 units from *the screwdriver*, 2.3 from *the wrench*).

In order to gain a more general view of the space as a whole, the semantic distances between all words of a particular pairing type in the similarity space (see Table 5 below) were calculated, to see whether words referring to objects and words referring to actions exhibit similarity effects that would distinguish between the two domains, despite the wide range of concepts included in each. These comparisons also included the set of nouns referring to actions, in order to test whether these words' representations follow their grammatical class (in which case they should be most similar to nouns referring to objects) or their semantic content (in which case they should be most similar to verbs referring to actions).

Table 5. Average word-word distances in semantic similarity space, within and across category. Distances are measured in arbitrary units based on ensemble averages of maps with dimensions 40x25 units (standard deviations in brackets)

	Grammatical class comparison		
	Noun-noun	Noun-verb	Verb-verb
Average distance	19.88 (4.05)	19.53(2.91)	16.12 (3.95)
	Semantic classification comparison (Within class)		
	Object noun	Action noun	Action verb
Average distance	18.90 (4.90)	15.94 (4.95)	16.12 (3.95)
	Semantic classification comparison (Between class)		
	Object N-action N	Object N-action V	Act.N-act. V
Average distance	21.91 (2.54)	20.70 (2.67)	16.53 (4.26)

First, the comparison between all nouns referring to objects and verbs referring to actions (169 objects and 216 actions) showed a distinction between these two groups of words. The average word-word distances (derived from ensemble average distances across the 100 self-organising maps with dimension 40x25, as described in Chapter 4) were 19.82 units within object-nouns (a total of 14,196 unique combinations of two object-nouns), 16.55 units within action-verbs (23,220 unique combinations of two action-verbs), and 20.35 units between action-verbs and object-nouns (36,504 possible pairings between one action-verb and one object-noun in the set). Action-verbs were more similar to other action-verbs than object-nouns were to other object-nouns; independent sample $\underline{t}$-tests comparing the 14,196 object-object pairings to the 23,220 action-action pairings revealed a significant difference ($\underline{t}(37414) = 57$, $\underline{p} < .0001$). This difference may have resulted from the differences in distances between semantic fields, with object-nouns tending to be more segregated into specific categories than action-verbs. The average distance between object-nouns and action-verbs (36,504 such pairs) was greater than either within-group average distance (within-object vs object-action

$t(50,698) = 21$, $p < .0001$; within-action vs. object-action $t(59,722) = 106$, $p < .0001$), demonstrating that the object-noun/action-verb distinction is reflected in the similarity space. This is a particularly important finding given the importance of this distinction in tasks involving meaning (Vigliocco, Vinson, Woolfe, Dye and Woll, 2005).

Investigating grammatical class: nouns referring to actions

Next, measures of semantic distances were used to calculate the extent to which grammatical class is reflected in FUSS. Grammatical class is one of the best candidates for being a language universal, and some researchers have claimed that grammatical class distinctions are emergent from semantic distinctions, beginning with the correspondence between objects and nouns on one hand, and actions and verbs on the other (e.g. Bates & MacWhinney, 1982, Elman, 2003). This would suggest that nouns and verbs should be semantically distinct from one another, a difference that should be reflected in behavioural tasks as well as neural organisation. Numerous studies in cognitive psychology, cognitive neuroscience and neuropsychology (reviewed in Vigliocco, Barber, Vinson, Druks & Cappa, in prep) have investigated this question by comparing performance of various impaired and unimpaired populations on tasks involving nouns and verbs, often showing substantial differences between the two. However, a majority of such studies fail to tease apart the grammatical class distinction between nouns and verbs from the conceptual distinction between objects and actions, so it is often unclear whether these findings are related to grammatical class or some other conceptual factor (see Vinson & Vigliocco, 2002; Vigliocco, Barber, Vinson, Druks & Cappa, in prep, for further discussion). In fact, the few studies of grammatical class that do control for conceptual factors have typically shown that grammatical class differences for nouns and verbs are not observed for tasks involving processing of (uninflected) single words (e.g. Chiarello, Liu, Shears & Kacinik, 2002; Vigliocco, Vinson, Arciuli & Barber, 2008; Vigliocco, Warren, Siri, Arcuili, Scott, & Wise, 2006).

This suggests that, contrary to views in which grammatical class is emergent from semantics, this might only be true for the strong correspondence between objects and nouns and not more generally of grammatical class. This question is examined through analysis of similarity within the lexical-semantic level of FUSS.

In order to do this, it was necessary to focus upon a set of words for which the (grammatical) distinction between nouns and verbs does not also correspond to the (semantic) distinction between objects and actions: nouns referring to actions such as *the blink* and *the scream*. The question asked was whether action-nouns are more similar to other object-nouns than to their action-verb counterparts (*to blink* and *to scream*). If grammatical class *per se* is an organising principle that goes beyond the semantic distinction between objects and actions, or if grammatical class is emergent from semantics, action-nouns should exhibit more similarity to object-nouns than to action-verbs. This could be expressed at various levels of specificity, ranging from the broadest (action-nouns should cluster among object-nouns and not among action-verbs) to the finest (action-nouns should exhibit a tendency to cluster among other action-nouns compared to similar action-verbs).

Considering all nouns together (object-nouns and action-nouns), the average within-grammatical class distances were 20.16 units for nouns (compared with only 16.55 for the verbs) . Moreover, the average distance between all noun-verb pairs was 19.49, a value significantly smaller than the average distance between pairs of object-nouns ($t(66{,}034) = 7.332$, $p < .0001$). This may be surprising until we consider the greater extent to which different categories of words referring to objects are distinguished from each other, while words referring to actions are much less separate (as illustrated by analyses of shared features and shared feature weights described in Chapter 3, and the analysis of distances between different semantic fields reported earlier in this chapter). By adding action-nouns to the set of object-nouns, the within-grammatical class coherence for the nouns was further reduced (and it was already fairly limited due to the categorical distinctions among object-nouns) and the difference between nouns and verbs was reduced. A follow-up distance analysis that compared distances within the class of action-nouns (17.40 units) to distances between action-nouns and verbs (17.33), and between action-nouns and object-nouns (20.39) showed that action-nouns were significantly closer to each other than to object-nouns ($\underline{t}(14{,}482) = 43$, $\underline{p} < .0001$), and were no closer to each other than to verbs despite the huge number of comparisons and resulting power to detect even small differences ($\underline{t}(17{,}819) = 0.554$, $\underline{p} = .579$). This shows that action-nouns are (semantically) more similar to action-verbs (in fact, indistinguishable from them in this type of analysis) than to object-nouns; that is, the semantic characteristics of actions vs. objects are responsible for the patterns of semantic similarity in FUSS, while grammatical class (verbs vs. nouns) does not have semantic consequences.

To further measure the extent to which action-nouns' representations are more similar to action-verbs than to object-nouns, distances of action-verb/ action-noun pairs (e.g. *to blink/the blink*, *to construct/the construction*) were examined. If there is correspondence between grammatical class and semantic similarity, action-nouns' representations should be closer to object-nouns in the similarity space than their minimal-paired action-verbs (e.g. *the blink/to blink*; *the construction/to construct*). To test this prediction, the average distance between a given action-noun and <u>all</u> object-nouns was calculated and compared to the distance between the corresponding action-verbs and all object-nouns. This measured the extent to which action-nouns are perceived as being closer to the centre of mass of the "object-noun space" than action-verbs. Because the measures of distance were so similar (reflecting the fact that action-nouns are represented very near their verb counterparts), non-parametric sign tests were used to test for the presence of any effect of grammatical class. Of the 71 action-noun/action-verb pairs, nouns were nearer to the object-nouns 36 times; and verbs was nearer 31 times (with 4 ties); a nonsignificant difference (sign test $\underline{p}$ = .625).

These results seem to indicate that grammatical class does not have semantic consequences, at least for this set of action-nouns. However, it may be possible that grammatical class effects are observed in a more subtle manner rather than being reflected in some kind of semantic properties common to all nouns. Properties of the semantic similarity space may preclude action-nouns from being represented near the object space because of intervening concepts, but some kind of nounlike characteristics within a semantic field might still be observable. In order to investigate this possibility, one more analysis of semantic distance was conducted to investigate whether effects of grammatical class could be observed within semantic fields referring to actions of different kinds. Here, only sets of action words within semantic fields were investigated. If grammatical class has consequences for semantic similarity, action-nouns should generally be closer to the action-noun member of a (different) action-noun/action-verb pair. For example, the action-noun *the request* should be closer to *the demand* than it is to *to demand*, *to demand* should be closer to *to request* than *the request*, and so on. In order to provide minimal semantic contrasts, only the following sets of words that unambiguously represented the same narrow semantic fields were selected: direction of motion (*ascent/ascend* vs. *descent/descend*); eye action (*blink*, *squint*, *wink*), noises (*clang*, *clatter*, *crackle*), light emission (*flicker*, *glow*, *shine*, *sparkle*), communication (*demand*, *plea/plead*, *request*, *suggest/suggestion*), vocal noises (*scream*, *screech*, *shout*, *yell*), facial expressions (*frown*, *smile*), body action (*pull*, *push*), and exchange (*trade*, *exchange*, *loan*). Within each field, distances were compared between each action-noun/action-verb pair and the other nouns and verbs only within each subset. Again non-parametric sign tests were used to test whether nouns tended to be closer to other nouns than to the corresponding verbs, and whether verbs tended to be closer to other verbs than to the corresponding nouns. Action-nouns were closer to other action-nouns in the same field in 15 comparisons, and to same-field verbs 14 times (1 tie), a non-significant difference ($p > .90$). A similar pattern appeared for verbs. Verbs were closer to the verb member of a pair 14 times, and closer to the noun 15 times. In short, there was no tendency for action-nouns to be semantically segregated from their action-verb counterparts: nouns and verbs appear to cluster together irrespective of grammatical class.

Taken together, the above results present a picture of the semantic similarity space obtained from the speaker-generated features which is highly consistent with intuitive judgements of semantic similarity, ranging from coarse-grained distinctions between objects and actions; to distinctions of moderate grain such as the segregation between different semantic fields: object-noun categories like fruit/vegetables, vehicles, animals, and action fields like cooking, light/heat emission, sounds, to fine-grained distinctions such as farm animals vs. wild animals vs. small mammals. They also reveal that the patterns of similarity among speaker-generated features are not related to grammatical class: action-nouns were not distinguished from action-verbs, and both were separable from object-nouns. This finding is in contrast with any account by which grammatical class has a conceptual or semantic basis.

These comparisons between distance measures in FUSS semantic similarity space and independently obtained semantic field relations reveal the utility of the dimensionality reduction techniques used in FUSS: similarity relations are respected, from very broad distinctions such as the divide between words referring to objects and to actions, down to very fine-level properties of similarity such as the semantic proximity of *apple, pear, peach*; or *get*, *receive, acquire*. However, a much more crucial test of the value of these measures is the extent to which they predict performance in behavioural tasks which are sensitive to semantic similarity, which will be addressed in the following chapters.

Chapter 6: Predicting fine-grained behavioural effects using measures of similarity from FUSS

The previous chapters have described properties of speaker-generated features and the development of FUSS, a model of lexical representation based upon them. Although the properties of the features in the conceptual level of FUSS and of the resulting semantic space are consistent with previous claims about semantic composition, and with intuitive notions of categories and similarity, a more direct test of FUSS is necessary. If the semantic similarity measures obtained from FUSS are indeed psychologically real, they should be able to make fine-grained predictions of semantic effects in behavioural tasks. That is, they should go beyond previous studies by showing that semantically-related words produce measurable effects compared to unrelated words (for example, the oft-replicated semantic priming effect), and more importantly, by predicting how such effects are modulated by the <u>degree</u> of semantic similarity.

Of crucial interest at this stage is the relative predictive power of the semantic distance measures in domains of objects and actions. In the previous chapters it became clear that both can be represented in this framework, but it remains to be seen whether these representational assumptions benefit words referring to objects and words referring to actions similarly. While graded semantic effects within the object domain are predicted by most models of semantic organisation, it is not clear whether this would be the case for the action domain, given that most models to date do not deal with this domain of knowledge. So in addition to testing for graded behavioural effects of the semantic distance measures in comprehension and production of words referring to objects, of critical interest will be the ability of the semantic distance measures to predict performance for words referring to actions.

Two complementary behavioural methodologies will be used to this end. One tests the effects of fine-grained semantic similarity in comprehension on the degree of semantic priming in a lexical decision task (Experiment 1: objects, Experiment 2: actions), and the second tests the degree of semantic interference in picture naming (Experiment 3: objects, Experiment 4: actions). Both sets of experiments will test, first, whether FUSS can predict graded semantic effects for words referring to objects, and, second, whether the same patterns of results are obtained for words referring to actions.

LEXICAL DECISION: SEMANTIC PRIMING

Semantic priming refers to the robust finding that speakers respond faster to a target word when preceded by a semantically related word than when it is preceded by an unrelated word (Meyer & Schvaneveldt, 1971; see Neely, 1991 for a review). The phenomenon of semantic priming has been extensively investigated because it arises in a largely automatic manner, and has been considered to reflect the organisation of semantic memory (e.g., Anderson, 1983; Collins & Loftus, 1975; Cree, McRae & McNorgan, 1999; McRae & Boisvert, 1998). Cree et al. (1999) and McRae and Boisvert (1998), have shown that categorical priming (i.e., words from the same semantic category such as jar-bottle; subway-bus; raft-canoe) can be observed even when those words are not associated, if the related items are selected on the basis of empirically obtained measures of semantic similarity, showing that such effects are not simply the product of word association reflected indirectly through speaker-generated features. Further, semantic priming effects appear to be symmetrical (e.g. *turkey* primes *goose* as much as *goose* primes *turkey; McRae & Boisvert, 1998*), consistent with the assumptions underlying the semantic distance model employed in FUSS, e.g., that a single distance measure between *turkey* and *goose* predicts behavioural effects in both directions.[11]

In this section two experiments are reported. Experiment 1 replicates previous work for words referring to objects, given measures of semantic distance from FUSS. Most importantly, however, this experiment goes beyond previous work, in that FUSS distances are not only used to distinguish between related and unrelated words, but also at a finer level, assessing the role of fine-grained degree of similarity among words that are somewhat related in meaning. Graded effects of meaning similarity are predicted under most theories of semantic representation, but to date there have been only a limited number studies of semantic priming which have explored this possibility using single words.[12]. Experiment 2 extends the investigation to the action domain, applying exactly the same methodology but using different items and participants. With respect to words referring to actions, most studies tend to investigate verbs in phrase or sentence contexts, and/or priming effects across grammatical classes (for example, *broom-sweep*) rather than priming from one action-verb to another (see Vigliocco, Vinson, Arciuli & Barber, 2008, for discussion) Only a few studies have reported semantic priming effects for verb-verb pairs (Rösler, Streb & Haan, 2001;Vigliocco et al., 2008; see also Bushell & Martin, 1997), and only using synonymous prime-target pairs (Bushell & Martin; Rösler et al.) or at least highly related pairs (Vigliocco et al.). The present study, therefore, will establish whether semantic priming effects for words

[11] Importantly, asymmetrical priming results can be observed when the relationship between words is primarily associative rather than semantic (e.g. Najmi & Wegner, 2008). It remains to be seen whether asymmetrical, purely semantic priming can be observed. If so, this would require a different set of processing assumptions than those employed here.

[12] Certain studies of morphological and phonological priming effects also embed items varying in semantic relatedness (in terms of the extent to which two words sharing the same stem or letter string are semantically related to each other), showing similar graded effects (e.g. Gonnerman, Seidenberg & Andersen, 2007). However, these studies typically focus upon gradation among words sharing orthographic or phonological overlap, with goals related to testing different account of morphology, while the present studies approach lexical-semantic representations across words typically without morphological or orthographic overlap.

referring to actions can be observed when prime and target are not synonyms, similarly to priming effects that have been repeatedly reported for nouns referring to objects.

Experiment 1: Objects

Method

Participants. Sixty-four native English speakers from the UCL community[13] participated in this study and received £3 for their participation. All participants reported having normal or corrected-to-normal vision. Nine participants who had high error rates (>10%) or extremely slow response times (> 3 SD's from other participants) were replaced.

Materials. Target items were selected from the list of words included in FUSS, and meeting various other restrictions as follows. First, target words and primes were all nouns depicting concrete objects, and were matched as closely as possible for verbal frequency, number of letters, and had minimal orthographic or phonological overlap with the target word. Primes were selected to be (1) very close to the target word (operationalised as word-word semantic distances in FUSS between 1.5 and 4.5), for example *dagger* – *sword*; (2) close (distances between 4.5 and 7.5), for example *dagger* – *razor*; (3) medium (distances between 7.5 to 10.5), for example, *dagger* - *hammer*; and (4) far (distances between 18 to 22), for example, *dagger* - *tongue*. Verbal frequency (Kucera & Francis, 1967) did not differ across conditions (average frequencies were 41.7 in the very close condition (SD = 16.7), 40.9 (14.6) for close, 41.5 (16.4) for medium and 42.0 (15.9) for far; ANOVA revealed that the four conditions did not significantly differ in

[13] Experiments 1-5 all included participants with various language backgrounds who reported their native language as English, and that they did not speak any other language fluently. Most of these participants spoke British English, but all experiments included some participants with other English backgrounds (Australia/New Zealand being most common, followed by US/Canada, and a few others such as South Africa and Singapore). Because these participants did not behave measurably differently from the British English speakers, their results were combined in all analyses reported.

terms of item frequencies: $\underline{F}(3, 93) < 0.1$), nor in terms of length in letters (average length = 4.97, 5.19, 5.09, and 4.97 respectively: $\underline{F}(3,93) < 0.1$). See Appendix B for a complete list of items used in the experiment.

Four lists were prepared from the experimental items, such that each target word appeared only once in a given list, and each prime word also appeared only once (although a single prime could appear for different targets across lists). Each list contained eight prime-target pairs from each of the four. There were altogether 32 experimental items in each list.

The experimental items were presented within the context of a large number of filler items including 40 noun-noun prime-target pairs (all selected to be unrelated at an intuitive level), and 72 verb-verb prime-target pairs. The large number of filler trials was selected based on previous semantic priming studies, as one way to avoid the possibility that priming effects may occur due to strategies (see Neely, 1991). An equal number of word prime - nonword targets were also included. Each nonword was created by taking a noun or verb not appearing elsewhere in the experiment and altering one letter, such that the resulting string was orthographically acceptable but was not a real word. Each word and nonword appeared no more than once for each participant. The resulting prime-target pairs were combined in a pseudorandom order such that each test item was separated from the next by at least two fillers.

Procedure. Participants were told that the experiment focused upon word identification processes, and that they would see a variety of words or nonwords, and their task was to indicate by button press whether a presented letter string was a word or not. "Word" responses were always made with the right hand, and "nonword" responses with the left. Participants were urged to respond as quickly as possible while trying to minimize errors. All item presentation and data collection used IBM PC-compatible computers running E-Prime software (Schneider et al., 2002). An initial practice set of 20 single-word trials ensured that the participants understood the lexical-decision task. The instructions then advised the participants that they would briefly see a word presented immediately before the target word or nonword, and they should attempt to ignore the prime word if possible, responding only to the target. Both the practice trials and experimental trials followed the same presentation mode: a central fixation point was displayed for 800 ms, followed by the prime word for 67 ms, followed by the target word, which remained on the screen until the participant pressed a button, followed by a 300ms blank interval. The short SOA between prime and target display (67 ms) was selected on the basis of previous studies showing that long SOAs permit strategic responding, and on the basis of a pilot study showing that a 67 ms SOA was sufficient to produce semantic priming (related vs. unrelated). The next trial began immediately thereafter. Every 100 trials the participant had the opportunity to take a short break. Reaction times and accuracy were recorded for each trial.

At the end of the experiment, participants were debriefed, with particular attention paid to whether they noticed any relationship between targets and primes. Those participants who noticed a relationship tended to focus upon the phonological similarity between filler targets and primes, or upon the wordlike properties of the nonword target items. No participants reported noticing any similarity in meaning between targets and primes.

<u>Design and data analysis</u>. The critical dependent measure was the time to respond, indicating that a critical target is indeed a word. Errors were recorded (false "nonword" responses to target words) and analysed separately. The independent variable was the semantic distance between target and prime, which was manipulated within (target) items and within subjects. The effects of semantic distance upon lexical decision reaction time were subjected to one-way ANOVA by subjects and items, with the linear trend component of particular interest. Trend analysis was performed using

contrast coefficients weighted on the basis of the semantic distances, both with subjects and items as random factors.

Results

Participants' responses to the critical items were highly accurate, an overall accuracy above 98%. Correct reaction times, averaged by condition are reported in Table 6 (standard errors of the mean in brackets) along with error frequencies.

Table 6. Average lexical decision latencies (RT, in ms; standard error of the mean in brackets) and error percentages as a function of semantic distance between target and prime. Experiment 1 (objects).

Semantic Distance	Response latencies	Error rate (%)
Very close	548 [8.1]	1.8 [1.2]
Close	557 [9.3]	1.5 [1.1]
Medium	567 [9.1]	1.8 [1.3]
Far	572 [9.8]	1.3 [1.1]

Reaction times. Reaction times for correct responses were collapsed by subjects and then by items, and subjected to a one-way analysis of variance. The reaction times were submitted to an omnibus analysis of variance, which was significant both by subjects and items ($F1(3,189) = 6.56$, $p < .001$; $F2(3,93) = 5.23$, $p < .001$) This was followed up by testing the linear trend using contrast coefficients [-1.3, -0.7, -0.1, 2.1] corresponding to the average distances between target and distracters [very close, close, medium, far], reflecting decreased priming as distance increased. This linear trend was significant both by subjects and items ($F1(1,61) = 5.21$, $p = .026$; $F2(1,31) = 4.48$, $p = .042$). Orthogonal quadratic and quintic trends were also tested: quadratic term, although significant by items, was only marginally significant by subjects: $F1(1,61) = 2.99$, $p = .089$, $F2(1,31) = 4.26$, $p = .048$; quintic term was nonsignificant either by subjects or by items $F1(1,61) = 1.79$, $p = .196$, $F2(1,31) = 1.62$, $p = .213$. These results indicate that priming effects were modulated by the semantic distance measures obtained from speaker-generated features.

Errors. Errors did not occur differently for primes from different semantic distances (all Fs<1).

Discussion

The main result from this experiment is the finding that the semantic distance measures modulated the amount of priming observed for words referring to objects. These results replicate and extend what has previously been reported by McRae and Boisvert (1998) and Cree et al. (1999) who also used speaker-generated feature norms to simulate priming effects within attractor network models. The novel finding here is that such an effect is modulated by the feature-based measure of semantic distance, a result consistent with the notion that fine-grained differences in similarity measures reflect gradations in semantic similarity. Next, a parallel experiment was conducted in the action domain to assess whether this is also true for this very different semantic domain.

Experiment 2: Actions

Method

Participants. Forty-eight native English speakers from the UCL community participated in this study and received £3 for their participation. All participants reported having normal or corrected-to-normal vision. Four participants who had high error rates (>10%) or extremely slow response times (> 3 SD's from other participants) were replaced.

Materials. Target and prime verbs were selected on the basis of the same distance criteria as Experiment 1. Verbal frequency did not differ across conditions (average frequencies from Kucera and Francis (1967) were 62.5 (SD = 28.5) in the very close condition, 52.9 (26.7) for close, 57.9 (25.4) for medium and 60.0 (28.1) for far;

ANOVA revealed these conditions did not significantly differ: F(3, 93) < 0.5), nor did length in letters (average length = 4.9, 4.9, 4.8 and 4.8 respectively, significantly different from each other: F(3,93) < 0.1). See Appendix C for a complete list of items used in the experiment.

Four experimental lists were prepared as in Experiment 1. The experimental items were presented within the context of a large number of filler items. In order to ensure that targets and primes were interpreted as verbs, filler items consisted of unambiguous verbs (not having a noun homonym) or words with verb-dominant frequency of use. There were 112 such verb-verb prime-target filler pairs (all intuitively "unrelated"). An equal number of word prime - nonword targets were included. Nonwords (based on verbs) were created in the same manner as in Experiment 1.The procedure, design and data analyses were also exactly the same as in Experiment 1.

Results

Participants performed the task with a high rate of accuracy (error rate for target items = 3.5%). Correct RTs were averaged across semantic distance by subjects and items; see Table 7 for response latencies and error rates.

Table 7. Average lexical decision latencies (RT, in ms; standard error of the mean in brackets) and error percentages as a function of semantic distance between target and prime. Experiment 2 (actions).

Semantic Distance	Response latencies	Error rate (%)
Very close	602 [8.6]	3.6 [1.9]
Close	613 [8.9]	3.1 [1.7]
Medium	627 [9.4]	4.0 [2.0]
Far	636 [10.0]	3.3 [1.8]

Reaction times. Reaction times for correct responses were collapsed by subjects and then by items, and subjected to a one-way omnibus analysis of variance. The effect of semantic distance was significant by subjects and items ($F1(3,141) = 5.01$, $p = .002$; $F2(3, 93) = 6.96$, $p < .001$) As in Experiment 1 this omnibus test was followed up by trend analysis, with the linear trend as the measure of interest. This trend was significant both by subjects and items ($F1(1,45) = 4.88$, $p = .032$; $F2(1,31) = 5.77$, $p = .023$); neither quadratic and quintic terms were significant: quadratic term: $F1(1,45) = 3.12$, $p = .084$; $F2(1,31) = 3.50$, $p = .071$; quintic term: $F1(1,45) = 1.94$, $p = .171$; $F2(1,31) = 3.53$, $p = .070$. Again, these results indicate that priming effects were linearly modulated by feature-based semantic distance measures.

Errors. Errors did not occur differently for primes from different semantic distances (all Fs<1).

Discussion

This experiment established that graded semantic priming in the action domain can also be observed, going beyond previous studies that have investigated verb-verb priming in which the effects of highly related primes are compared to unrelated primes (e.g., Bushell & Martin, 1997; Rösler et al., 2001; Vigliocco et al., 2008). Hence, this experiment provides evidence in a lexical decision task, from the domain of actions, that important aspects of similarity among lexical-semantic representations can be captured across domains using the same general computational principles. In order to test the generality of these effects, it is important that they can also be observed in another behavioural domain in which semantic effects have been reported: the picture-word interference paradigm.

PICTURE NAMING LATENCIES: SEMANTIC INTERFERENCE

In contrast to the facilitatory semantic effects arising in primed lexical decision (as reported in Experiments 1 and 2), semantically related words exert interfering effects during picture naming, as shown in picture-word interference experiments in which a distracter word is presented immediately before a target picture to be named. In these experiments, speakers are slower to name the picture when the word is semantically related to the target than when the word is unrelated (Glaser & Düngelhoff, 1984; Schriefers et al., 1990)[14]. Interference effects have been reported for both object-nouns (e.g., Glaser & Düngelhoff, 1984; Lupker; 1979; Schriefers et al., 1990, and many others) and action-verbs (Roelofs, 1993; Vigliocco, Vinson & Siri, 2005). These studies, however, only contrasted related and unrelated words.. The experiments reported below extend previous work by manipulating the degree of semantic relatedness between the distracter word and the target picture name on the basis of the semantic distance measures in FUSS in the same way as in Experiments 1 and 2 above, to investigate whether the degree of similarity affects the amount of interference in naming, In Experiment 3, participants are asked to name pictures of objects while ignoring distracters which are object-nouns varying in semantic distance to the target noun, and in Experiment 4 participants name pictures of actions while ignoring distracter words referring to actions. If semantic distance predicts performance in this task, semantically related distracters should affect naming latencies as a function of their distance to the target word.

Experiment 3: Objects

Method

Participants. Thirty-six native English speakers from the UCL community

[14] The difference between the direction of these semantic effects (facilitation in lexical decision, interference in picture naming) can be explained in terms of differences between the tasks. In lexical decision, participants are only required to recognise whether a given string of letters is a word or not, while in picture naming they must select and articulate a specific word, without any orthographic information being present (because the input is a picture). In this latter case, other semantically-related lexical representations could slow down the selection/naming process by competing to be selected. For lexical decision, no such competition arises because it is not necessary for a unique word to be selected.

participated in exchange for monetary compensation. All participants reported having normal or corrected-to-normal vision. Six participants whose responses were unsuitable (e.g. stuttering or speaking too quietly to trigger the voice relay) were replaced.

Materials. Twenty-four target pictures were selected, along with distracter words. Target pictures were the pictures labelled by a subset of the object-nouns from the feature set. The nouns that were included among the targets all had high levels of name agreement (as indicated by ratings in Snodgrass and Vanderwart, 1980, and confirmed to hold for speakers of UK English based in a pilot study). Distracter words were selected on the basis of semantic distance as in Experiments 1 and 2: very close (1.5 to 4.5 units in FUSS lexical-semantic distances), close (4.5 to 7.5 units), medium (7.5 to 10.5 units), and far (or unrelated, distance > 18.5 units). Distracters never appeared as targets, and targets never acted as distracters. Distracters were also phonologically dissimilar to their targets.

Distracters were matched as closely as possible for frequency and length. Average frequency (Kucera & Francis, 1967) for distracters in the very close, close, medium and far distances were respectively (42.5 (SD=17.3), 42.9 (SD=15.6), 42.1 (SD=15.9) and 42.7 (SD=14.9)). Repeated measures ANOVA revealed no significant differences between these groups, $F(3,69) = .21$, $p = .60$. Word length also did not significantly differ between conditions (mean length was 5.21, 5.38, 5.33, and 5.13 letters); $F(3,69) = .109$, $p = .744$. A full list of materials used in this experiment can be found in Appendix D. Twenty-four filler pictures were also selected from semantic fields distinct from those represented in the target items (e.g. miscellaneous artefacts, plants, musical instruments); four distracter words for each filler picture were selected, one of which was from the same semantic field as the target picture, and three of which were from different semantic fields.

The experimental structure consisted of four blocks of 48 trials each (24 targets, 24 fillers). Target pictures were divided into four sets, so that an equal number of target-distracter pairs from each semantic distance would appear in each block, and each target picture would appear only once in each block. For example, the six target pictures from Set A might appear in the first block with very close distracter words, in the second block with medium distracters, in the third with close distracters, and in the final block with far distracters. Different sequences of blocks were assigned to four different lists using an incomplete Latin square design. For the purpose of preparing lists, each filler picture was arbitrarily paired with a target picture and assigned to blocks in a parallel manner. Items in a block within each list were presented to each participant in a pseudorandom order, with the only constraint being that target pictures and filler pictures alternated.

Procedure. Participants were told that the experiment investigated word production processes. They were asked to name the pictures as quickly as possible trying to be as accurate as possible and ignoring distracter words. Item presentation and data collection used IBM PC-compatible computers running E-Prime software (Schneider et al., 2002). Vocal response latencies were measured using an E-Prime Deluxe Serial Response Box. Responses were also tape-recorded and monitored online for accuracy.

An untimed picture naming phase started the experiment. The pictures were presented to the participants, who were asked to name the picture aloud. This ensured that they recognised the pictures and confirmed that the pictures had high name agreement. A sequence of practice trials followed this phase. During the practice trials, the target pictures and filler items were presented once paired with an unrelated distracter. If a participant exceeded 10% errors, another practice block was performed.

Once the practice trials were complete, the experimental trials, divided into four blocks, were administered. A fixation cross appeared at the centre of the screen for 500 milliseconds, followed by a 50 millisecond blank screen. The distracter word then appeared on the screen in a randomly-selected location either above or below the fixation cross, and 150 milliseconds later the target picture appeared in the location not occupied by the distracter. The target picture and distracter remained on the screen until triggered by vocal response (i.e. picture naming), or 2500 milliseconds elapsed. A blank screen was displayed for 300 milliseconds, followed by the fixation cross for the next trial. Between blocks participants were given the opportunity to take a short break.

Design and data analysis. The independent variable (manipulated within subjects and items) is the semantic distance between target and distracter. The main dependent measure was the duration between presentation of a target picture and a participant's (correct) response as measured by a voice relay. A secondary dependent measure was the number of errors. One-way analysis of variance (with trend contrast coefficients) was carried out upon the difference scores, with particular attention to the linear trend. As in the previous experiments reported here, trend contrast coefficients were weighted on the basis of the semantic distances, both with subjects and items as random factors.

Three types of errors were identified. Content errors were scored when the participant made an error in naming the target picture (including semantic substitution errors, recognition errors, stutters and dysfluencies). Detection errors were scored when the voice relay failed to detect the correct word onset (voice relay triggered too early, too late, or not at all). Finally, trials with response latencies deviating more than three standard deviations (by participant and condition) were classified as outlier errors. Analyses of variance, using semantic distance as the independent variable, were conducted on each error category separately.

Results

Naming latencies. All three types of errors were excluded from the naming latencies analyses. Correct response latencies are reported in Table 8.

Table 8. Average picture naming latencies (RT, in ms; standard error of the mean in brackets) and error percentages of different types as a function of semantic distance between target and distracter. Experiment 3 (objects).

		Error type (%)		
Semantic Distance	Response latencies	Content	Detection	Outlier
Very close	671 [8.8]	5.4	1.2	0.9
Close	657 [8.2]	4.4	1.4	1.3
Medium	648 [7.9]	4.2	0.9	1.1
Far	642 [8.0]	3.6	1.1	1.2

First, omnibus ANOVA was performed to test for the general effect of semantic distance. This was significant by subjects and items ($F1(3,105) = 4.27$, $p = .007$; $F2(3,69) = 5.11$, $p = .003$). Linear trend analysis was then performed using within-subjects and within-items ANOVA, using contrast coefficients [-1.3, -0.7, -0.1, 2.1] corresponding to the average distances between target and distracters [very close, close, medium, far] to test the hypothesis of linear trend. This linear trend component was significant both by subjects and items ($F1(1,35) = 5.59$, $p = .024$, $F2(1,23) = 4.71$, $p = .041$), indicating that the semantic interference effect was modulated by semantic distance measures obtained from the speaker-generated features. Tests of orthogonal quadratic and quintic trends revealed only marginal significance of either: quadratic: $F1(1,35) = 3.30$, $p = .078$, $F2(1,23) = 2.58$, $p = .122$; quintic: $F1(1,35) = 2.66$, $p = .120$, $F2(1,23) = 1.99$, $p = .172$, indicating that the primary effect of semantic distance was linear.

While very close and close distracter-picture pairs were always from the same semantic field, and far pairs never were, word-picture pairs in the medium distances range were not always from the same category. Distracter words at medium distance were sometimes from the same category as the picture and at other times from a different category, because otherwise it would have been impossible to have a sufficient number of items.. To check whether naming latencies were affected by this, a post-hoc test was conducted. Among the distracters of medium distance, 6/24 were from the same category as the target. Comparing the interference effects for these six items to the 18 items from different categories, no significant difference was observed ($t<1$). In other words, there was no additional benefit of category membership beyond feature-based semantic distance for these items.

Errors. Analyses of variance by subjects and items were performed to determine whether errors (overall error rate reported in Table 8) of different types systematically varied by semantic distance between target and distracter. Only content errors were affected by semantic distance, $F1(3,69) = 6.22$, $p < .01$, $F2(3,141) = 5.91$, $p = .02$ (errors were more common for distracters from closer semantic distances), all other Fs < 1.

Discussion

The finding of an overall interference effect is not surprising, as it has been observed in a number of previous studies; as in Experiments 1 and 2 the novel finding is the modulation of this effect by FUSS semantic distance measures: interference was greatest for the most similar distracters according to FUSS measures. Next, a parallel experiment was conducted in the action-verb domain to assess whether FUSS measures of semantic similarity are equally good at predicting performance in this domain in terms of graded interference effects.

Experiment 4: Actions

Method

Participants. Forty-eight native English speakers from the UCL community participated in exchange for monetary compensation. All participants reported having normal or corrected-to-normal vision. Eight participants' responses were inconsistent (e.g. stuttering or speaking too quietly to trigger the voice relay), and one participant was unable to comprehend the task; these participants were replaced.

Materials. Twenty-one target pictures depicting actions with high levels of name agreement (Druks & Masterson, 2000) were selected, along with associated distracter words meeting a series of criteria similar to those of Experiment 1. A critical additional criterion both for targets and distracters was that their names, if homonymous with nouns, were required to have a dominant frequency of verb usage (>90%). Distracter words were selected from the set on the basis of semantic distance: very close (1.5 to 4.5 units of FUSS lexical-semantic distance), medium (4.5 to 10.5 units, but favouring items closer than 7.5 units whenever possible), and far (or unrelated, distance > 18.5 units). Distracters were matched as closely as possible for frequency, length, and to minimise phonological dissimilarity to the target word. Frequencies (Kucera & Francis, 1967) did not significantly differ between conditions (mean frequency was 58.2 (SD=26.3) for very close, 57 (24.6) for medium, and 58.7 (25.7) for far distance; $F(2, 40) < 0.2$), nor did length (mean lengths 4.57, 4.67, 4.62 respectively, $F(2,41) < 0.1$). Only targets for which suitable distracters could be found at each distance were included; also, distracters never appeared as targets, and vice versa. A full list of the items used in this experiment can be found in Appendix E. One filler for each experimental item was created as in Experiment 3. Twenty-one filler pictures were selected; and three distracter words for each filler picture were selected, one of which was similar to the target picture (intuitively judged), and two of which were unrelated.

The experimental structure consisted of three blocks of 42 trials each (21 targets, 21 fillers). Target pictures were divided into three sets, so that an equal number of target-distracter pairs from each semantic distance would appear in each block, and each target picture would appear only once in each block (as in Experiment 3). Different sequences of blocks were assigned to three different lists using an incomplete Latin square design. Blocks were otherwise treated the same as in Experiment 3.

Procedure. The same basic procedure was followed as in Experiment 3. Participants were instructed to name each action using a stem+"ing" (e.g. *jumping*, *walking*), a response type heavily favoured by participants naming action pictures in unconstrained settings. Design and data analyses were the same as in Experiment 3, with the exception that there were only three semantic distance conditions in the present Experiment.

Results

Naming latencies. As in Experiment 3, all trials in which an error was recorded, or in which response latencies deviated more than three standard deviations (by participant and condition) were excluded from the naming latencies analyses. Correct response latencies are reported in Table 9.

Table 9. Average picture naming latencies (RT, in ms; standard error of the mean in brackets) and error percentages as a function of semantic distance between target and distracter. Experiment 4 (actions).

		Error type (%)		
Semantic Distance	Response latencies	Content	Detection	Outlier
Very close	790 [8.0]	6.4	0.9	0.7
Close-medium	778 [6.3]	4.4	1.1	0.9
Far	761 [7.6]	4.4	0.8	0.9

First, omnibus ANOVA assessed the effect of semantic distance, which was significant by subjects and items ($\underline{F1}(2,94) = 9.42$, $\underline{p} < .001$; $\underline{F2}(2,40) = 7.90$, $\underline{p} < .001$) Trend analysis was performed using within-subjects and within-items ANOVA, using contrast coefficients [-0.7, -0.3, 1.0] corresponding to the average distances between target and distracters [very close, medium, far] to test the hypothesis of linear trend. The linear trend component was significant by subjects and items ($\underline{F1}(1,47) = 11.68$, $\underline{p} = .001$, $\underline{F2}(1,20) = 8.91$, $\underline{p} = .007$), while the corresponding orthogonal quadratic component was not significant (both $\underline{F}$s < 1.2).

Errors. Error frequencies are reported in Table 9. Analysis of variance by subjects and items was performed to determine whether errors of different types systematically varied by semantic distance between target and distracter. Only content errors were affected by semantic distance ($\underline{F1}(2,61) = 3.21$, $\underline{p} = .046$; $\underline{F2}(2,142) = 4.00$, $\underline{p} = .028$; all other $\underline{F}$s (by subjects and items) < 1).

Discussion

The main result from this experiment, beyond replicating the finding of semantic interference for verb naming (Roelofs, 1993; Vigliocco, Vinson & Siri, 2005), is the observation of a modulation of the interference effect for words referring to actions, parallel to the effect observed for the object-nouns. Thus, the results of Experiments 3 and 4 converge in indicating that parallel effects can be observed for object and action domains, in support of the notion that a common semantic distance model based on speaker-generated features can predict performance in both domains despite the various differences in featural composition described in Chapter 3.

General Discussion

In all four experiments, FUSS semantic distance measures predicted fine-grained performance on tasks sensitive to semantic similarity, both for object and action domains. These experiments provide important evidence that the semantic representations of words referring to objects and words referring to actions can be based on the same general principles despite the numerous differences between the content and functions of such words, and also that important aspects of the semantic organisation of words referring to actions can be captured through the assumptions of FUSS. Most importantly, graded effects were observed across the four experiments. Although these findings are consistent with many models of semantic representation, this is the first time such predictions have been tested in semantically complex domains of knowledge. Graded effects have previously been reported only in content domains of colour and number (Klopfer, 1996; Brysbaert, 1995; Moyer & Landauer, 1967; Pavesi & Umiltá, 1998) which may be special content domains for which it is easy to describe the underlying conceptual dimensions (hue and saturation for colours, quantity for numbers). This finding of gradation is particularly important in the action domain which has received much less attention in behavioural studies of this type, showing that gradation in similarity among representations of words referring to actions also is a good predictor of semantic effects despite the many differences between object and action domains (see Vigliocco et al., 2004).

Although most other models of semantic representation also predict graded effects (at least for nouns referring to objects), whether on the basis of shared features, length of network connections, proportion of shared hidden units or proximity in attractor space, there are some exceptions where at least some knowledge (e.g. evolutionarily distinct categories) is strictly categorical in nature (e.g. Caramazza & Shelton, 1998). In such cases, these categories should be fully distinguished from each other and should not exhibit gradation. We will return to this issue in Chapter 8. In the next chapter, the behavioural results reported here are used to test the relative ability of

FUSS to predict fine-grained semantic effects for object and action domains, compared to other models which also predict graded semantic effects.

Chapter 7: Comparing models of fine-grained semantic effects

The previous chapter demonstrated that fine-grained semantic effects can be predicted by FUSS, a model of lexical-semantic representation based on speaker-generated features. The next question is whether other models of word meaning can also predict performance to a similar degree. After all, at least considering the object domain, it seems that nearly all models of semantic organisation would predict graded semantic effects (McRae; HAL; LSA; Wordnet; Network Models[15]; Semantic Fields). It is not clear, however, whether this would be the case for the action domain given that models usually do not discuss this domain of knowledge.

Moreover, with the exception of two types of models: global co-occurrence models such as LSA (Landauer & Dumais, 1997) and HAL (Burgess & Lund, 1997) and certain hierarchical network models such as Wordnet (Aguirre & Rigau, 1996; Budanitsky & Hirst, 2001; Fellbaum, 1990; Miller, 1995; Miller & Fellbaum, 1991; Richardson, Smeaton, & Murphy, 1994), existing models of semantic organisation do not allow us to empirically evaluate graded effects. For example, connectionist models designed to account for semantic priming often use artificially-generated semantic representations instead of real words (e.g. Plaut, 1995 used "category prototypes", which were each a random pattern of activation across 100 semantic features, and generated "category exemplars" by randomly altering some of the features of a given prototype in a designated manner). Other models that have used actual words, such as McRae et al. (1997) and other similar models, are limited to the object domain. The models that allow the derivation of quantitative predictions, however, differ among themselves in terms of goals, representational assumptions and implementations. Comparing the predictive power of similarity measures derived from FUSS to those derived from LSA and Wordnet allows an evaluation of the assumptions on which the different models are based. Here, using the data from Experiments 1-4 reported in the previous chapter, FUSS semantic distances are formally compared to similarity measures obtained from LSA and Wordnet.

Method

Operationalisation of Similarity Measures

In order to draw comparisons between the different models' ability to predict performance, it was necessary to obtain measures of semantic similarity between pairs of words for each. For FUSS, these are the distances described in previous chapters. For LSA, measures of semantic similarity (cosines between words' representations in similarity space, higher values reflecting greater proximity) were obtained through LSA's web-based interface (http://lsa.colorado.edu), using the "General reading up to 1st year of college" topic space and the "Matrix comparison" application. For Wordnet, measures were obtained using the Wordnet 1.6 database (Miller & Fellbaum, 1991). Wordnet has a hierarchical link structure between representations, and a measure of semantic distance between two words was obtained by counting the number of hypernym/hyponym (superordinate/subordinate) links between them, based on the nearest shared hypernym (e.g. the most-specific shared superordinate term), using software developed by Lewis (2002). In Wordnet, homonyms and polysemous forms are encoded with different senses, therefore it was necessary to sense-encode each target word explicitly. This was done with reference to the target pictures used in the

[15] Network models only predict graded effects if they contain more information than simply <ISA> links between subordinates and superordinates. The latter type of strictly hierarchical models would only predict graded effects between categories (depending on shared superordinates at higher levels) and no graded effects among members of the same category.

experiment: senses best corresponding to each target picture were selected. In the case of ambiguous coding (e.g. *carrot* has two similar senses: "deep orange root of the cultivated carrot plant" and "orange root, important source of carotene"), distances based on both senses were evaluated; one set of distances was selected for such items, based on proximity to other exemplars from that semantic field. For the *carrot* example, the second sense was more similar to that of other vegetables (because it falls into "food" hierarchy, while the first sense falls into "plant organ" hierarchy), so it was selected for analysis.

Additional assumptions were necessary in order to obtain Wordnet distances for verbs referring to actions. This is a consequence of the differing organisation of the object and action spaces within Wordnet. Whereas object nouns are organised hierarchically on the basis of hypernym/hyponym relations, the action spaces are instead organised in a more complex manner, including hypernym/troponym (manner of doing something), entailment and antonymy. A consequence of this organisation is that finding the shortest path between two action verbs in Wordnet is more complex than doing the same within the object-noun space where it is almost always solved by identifying a common hypernym. Again, software designed by Lewis (2002), configured to investigate hypernym/ troponym/ entailment/ antonymy/ synonymy links, where the length of the shortest path indicates the degree of similarity, was used.. A further complication is that many of the target verbs have more than one sense (and also have more senses than the object nouns). For example, the verb *run* has 42 distinct senses. This renders the sense-coding effort more difficult. Nevertheless, the sense-coding criteria used for the object nouns remained useful in identifying the sense that most closely corresponds to a pictured action. Finally, in Wordnet there are multiple separate clusters of action verbs ("verb files") which are represented independently (i.e., no links exist between them). For verbs in separate clusters, a high value was assigned to the

Wordnet distance measure (n=12 nodes), comparable to the largest distance found in the experimental set of items to indicate their relative degree of isolation under Wordnet's similarity structure.

Design and Analysis

In order to compare FUSS to LSA and Wordnet semantic similarity measures, multiple regression models were used to establish the predictive power of each similarity measure on lexical decision RTs or naming latencies. Because all experiments were conducted within (target) items, multiple regression was performed, with target items treated as a categorical factor using dummy variables, entered at the first step, the residuals of which were passed to a second step in which one of the semantic similarity measures was used as a predictor. The resulting regression model fit separate parallel lines for each target item, whose slope corresponds to the effect of semantic similarity controlling for item variability. Because the semantic similarity measures tend to correlate with each other, partial correlations between the dependent measure and a given semantic similarity measure (controlling for variation due to items) were compared using Meng, Rosenthal and Rubin's (1992) Z test. These analyses were conducted to compare the level of correlation between RTs (or naming latencies) and FUSS distances, between RTs and LSA, and between RTs and Wordnet similarity measures.

Sequential multiple regression was performed in order to contrast the performance of the different semantic similarity measures. With lexical decision RT or naming latency as a response variable, a first step was always to enter in predictors consisting of dummy variables to code for target item, thus removing target item-specific variability from the RT data. In a second step, either FUSS distances, LSA similarity measures, or Wordnet node counts were entered.

Results

Experiment 1: Semantic Priming for Objects

For the model including FUSS, model R^2 = .925 (regression ANOVA F = 36.845, p<.001); partial correlation between FUSS and RTs controlling for items was .665 (feature distance term, t = 8.683, p<.001). For the model including LSA, model R^2 = .893 (regression ANOVA F = 24.482, p<.001); partial correlation between LSA and RTs controlling for items was -.442 (LSA term, t = 4.782, p<.001). For the model including Wordnet, model R^2 = .902 (regression ANOVA F = 27.421, p<.001); partial correlation between Wordnet and RTs controlling for items was .519 (Wordnet term, t = 5.922, p<.001).

Partial correlations between semantic predictors and naming latencies (controlled for item variability) were compared using Meng et al.'s (1992) Z test on correlation magnitudes. For the comparison between FUSS and LSA (measures correlated at r = -.5202 for this set of items), Z = 3.269, p < .001 (one-tailed), indicating that FUSS was a better predictor of lexical decision RTs than was LSA. For the comparison between FUSS and Wordnet (measures correlated at r = +.7313), Z = 2.884, p = .002, indicating that FUSS was also better than Wordnet at predicting RTs. LSA and Wordnet did not significantly differ on this measure (p = .3208).

Experiment 2: Semantic Priming for Actions

For the model including FUSS, model R^2 = .877 (regression ANOVA F = 21.084, p<.001); partial correlation between FUSS and RTs controlling for items was .699 (feature distance term, t = 9.535, p<.001). For the model including LSA, model R^2 = .824 (regression ANOVA F = 13.923, p<.001); partial correlation between LSA and RTs controlling for items was -.589 (LSA term, t = 6.876, p<.001). For the model including Wordnet, model $\underline{R}^2$ = .778 (regression ANOVA $\underline{F}$ = 10.416, $\underline{p}$<.001); partial correlation between Wordnet and RTs controlling for items was .286 (Wordnet term, $\underline{t}$ = 2.909, $\underline{p}$=.005).

Partial correlations between semantic predictors and naming latencies (controlled for item variability) were compared using Meng et al.'s (1992) $\underline{Z}$ test on correlation magnitudes. For the comparison between FUSS and LSA (measures correlated at $\underline{r}$ = -.4821 for this set of items), $\underline{Z}$ = 1.758, $\underline{p}$ =.039 (one-tailed), indicating that FUSS was better at predicting naming latencies than was LSA. For the comparison between FUSS and Wordnet (measures correlated at $\underline{r}$ = +.2774), $\underline{Z}$ = 4.854 $\underline{p}$ < .001, indicating that FUSS was also better than Wordnet at predicting latencies. LSA similarity was also significantly better than Wordnet ($\underline{Z}$ = 3.245; $\underline{p}$ < .001).

<u>Experiment 3: Picture-Word Interference for Objects</u>

For the model including FUSS, model $\underline{R}^2$ = .891 (regression ANOVA $\underline{F}$ = 24.267, $\underline{p}$<.001); partial correlation between distances and RTs controlling for items was -.640 (feature distance term, $\underline{t}$ = 7.305, $\underline{p}$<.001). For the model including LSA, model $\underline{R}^2$ = .854 (regression ANOVA $\underline{F}$ = 17.115, $\underline{p}$<.001); partial correlation between LSA and RTs controlling for items was .455 (LSA term, $\underline{t}$ = 4.452, $\underline{p}$<.001). For the model including Wordnet, model $\underline{R}^2$ = .871 (regression ANOVA $\underline{F}$ = 19.948, $\underline{p}$<.001); partial correlation between Wordnet and RTs controlling for items was -.546 (Wordnet term, $\underline{t}$ = 5.717, $\underline{p}$<.001).

Partial correlations between semantic predictors and naming latencies (controlled for item variability) were compared using Meng et al.'s (1992) $\underline{Z}$ test on correlation magnitudes. For the comparison between FUSS and LSA (measures correlated at $\underline{r}$ = -.5493 for this set of items), $\underline{Z}$ = 2.480, $\underline{p}$ =.0066 (one-tailed), indicating that FUSS was a better predictor of naming latencies than was LSA. For the comparison between FUSS and Wordnet (measures correlated at r = +.7267), Z = 1.653, p = .0492, indicating that FUSS was also better than Wordnet at predicting latencies. LSA similarity measures and Wordnet did not significantly differ on this measure (p = .2857).

Experiment 4: Picture-Word Interference for Events

For the model including FUSS, model R^2 = .955 (regression ANOVA F = 41.792, p<.001); partial correlation between feature-based distances and RTs controlling for items was -.681 (feature-based distance term, t = 5.961, p<.001). For the model including LSA, model R^2 = .945 (regression ANOVA F = 32.671, p<.001); partial correlation between LSA and RTs controlling for items was .577 (LSA term, t = 4.464, p<.001). For the model including Wordnet, model R^2 = .921 (regression ANOVA F = 22.885, p<.001); partial correlation between Wordnet and RTs controlling for items was -.238, a nonsignificant correlation (Wordnet term, t = 1.567, p=.125).

Partial correlations between semantic predictors and naming latencies (controlled for item variability) were compared using Meng et al.'s (1992) Z test on correlation magnitudes. For the comparison between FUSS and LSA (measures correlated at r = -.5790 for this set of items), Z = 1.221, p =.111 (one-tailed), indicating that FUSS was not significantly better at predicting naming latencies than was LSA. For the comparison between FUSS and Wordnet (measures correlated at r = +.3754), Z = 3.717, p < .001, indicating that FUSS was better than Wordnet at predicting latencies. LSA similarity was also significantly better than Wordnet (Z = 2.211; p = .027).

Discussion

The results of the model comparisons show not only that FUSS consistently predicts the degree of semantic effects observed in primed lexical decision and picture-word interference (as already illustrated to some extent by the linear contrasts described in the Results of Experiments 1-4), but also that FUSS's predictive power is superior to that of the other models tested. FUSS measures were superior to LSA-based measures for three of the four experiments, and superior to Wordnet-based measures for all four.

The differences in performance between these models might be attributed to differences in the models themselves. LSA is primarily focused upon issues of acquisition and the development of semantic representations from a given input (particularly, extraction of meaning relations from text). Because of this, it is not necessary that LSA's representations are interpretable in any manner beyond abstractly representing a word's meaning in the context of other words (Burgess & Lund, 1998). The present focus, instead, is directly upon meaning representation, thus the information from which the FUSS similarity space is developed must be interpretable in order to allow us to evaluate assumptions concerning featural representations. They are thus constrained by neuroanatomical considerations grounding the featural descriptions in a manner that is not possible (or even desirable) for models such as LSA and HAL (Glenberg & Robertson, 2000). The greater predictive power of FUSS over LSA may be plausibly related to the different focus of the two approaches: because LSA is concerned with extracting meaning information from text, it cannot avoid embedding a certain degree of noise due to to homonymy and polysemy. Given the present focus upon meaning representation, words were selected and speaker-generated features were gathered in a manner designed to avoid homonymy/ polysemy as much as possible.

FUSS also outperformed similarity measures derived from Wordnet in all of the experiments. With respect to words referring to objects (Experiments 1 and 3), the poorer performance of Wordnet-based similarity may be a consequence of its network structure. Nouns are organised hierarchically, which has a very straightforward consequence: all words under the same mother node and linked to the mother node by the same relational link are equidistant. This implies the impossibility of graded effects between any pair of words under a given mother node. This has less of a detrimental effect in the object domain because of the large number of hierarchical levels in the object representation space. However, the same is not true for verbs referring to actions; Wordnet-based similarity was a very poor predictor of behavioural results for the action domain (Experiments 2 and 4). In particular, this measure was worse than FUSS or LSA in both experiments and did not even reach significance in Experiment 4. Unlike objects, actions in Wordnet are organised into far fewer levels, and into isolated networks. For example, intuitively *cough* and *spit* are very similar to *sneeze*. However, within Wordnet, *cough* and *spit* are represented within a network of words referring to acts of expulsion, while instead *sneeze* is represented in an independent network referring to involuntary acts. Both the existence of isolated networks and lack of depth in the hierarchical organisation within each network may contribute to Wordnet's poor predictive performance in the domain of actions.

It is important to consider, however, that some of the success of FUSS at predicting the semantic effects in Experiments 1-4 may be related to the fact that these experiments were designed in a manner that could have favoured FUSS over Wordnet or LSA. After all, all of the items in these experiments passed through selection and pre-processing before they were included in FUSS, and they were selected on the basis of FUSS distances. Had these distances failed to correspond to some extent with intuition about semantic relatedness, these experiments would probably not have been carried out. The same is not true of Wordnet or LSA; the items were selected and the experiments carried out before these models were consulted. There is therefore an element of circularity involved; an ideal basis for comparison would involve a set of items chosen without reference to FUSS, Wordnet or LSA. Unfortunately FUSS has a very limited vocabulary for this purpose (most of the items used in published semantic priming studies are not included in FUSS), and extending FUSS to include a substantially greater number of words would be an extremely time-consuming effort.

In general, however, the results of these model comparisons show that FUSS can not only predict the degree of semantic effects for both object and action domains, but also that its predictive power is superior to other extant models of representation that allow the extraction of item-specific similarity measures. These results, however, are still relatively constrained, referring only to fine-grained similarity among pairs of words that are reasonably closely related. However, FUSS also makes predictions at a relatively coarser level--the relative proximity between semantic categories or groups of words (e.g. the analyses of between-field semantic distances reported in Chapter 6). In the next chapter, Experiment 5 tests whether proximity at this level also has behavioural consequences.

Chapter 8: Testing category-level predictions of FUSS

The results of the four experiments reported in Chapter 6 clearly demonstrated that FUSS was a strong predictor of behavioural performance at the item-level: the degree of semantic priming or interference between two words was highly predicted by the feature-based distance between them in the object and action domains. Coarser-grained properties of similarity, measured as proximity among categories, were also observed, as discussed in the analyses of category-level semantic distances in Chapter 5. If it could be shown that these latter patterns of proximity are also reflected in behavioural performance, it would provide additional evidence for the implementational assumptions underlying the development of FUSS.

Experiment 5: Semantic blocking in picture naming

The "semantic blocking effect" in picture naming (Damian, Vigliocco & Levelt, 2001; Kroll & Stewart, 1994) arises when speakers are asked to name pictures in the context of other pictures. When the pictures in a given block are from the same semantic field, naming a picture is slower than for the same picture when it is presented in a block with semantically-unrelated pictures. It is generally agreed upon that such effects arise from semantic competition during the conceptually-driven lexical retrieval process (Levelt, Roelofs & Meyer, 1999).

Here, the semantic blocking effect is used to test whether graded similarity effects among groups of items can be observed, and whether these effects are similar for words referring to objects and words referring to actions. Within the object domain, as described previously, category membership has powerful effects, most striking in patients who are selectively impaired or spared in one category of knowledge, such as animals (Caramazza & Shelton, 1998); body-parts (Shelton, Fouch & Caramazza, 1998) and fruits and vegetables (Hart, Berndt & Caramazza, 1985). These findings have led some researchers to postulate that domains playing a fundamental role for our survival (e.g. animals, plants and body-parts) are represented categorically in semantic memory within dedicated neural substrates (Caramazza & Shelton, 1998). In this view, semantic distance effects may not be observed between evolutionarily motivated categories. These should act as isolable clusters because they are independent from other domains of knowledge. In contrast, graded effects may be observed between categories which are not evolutionarily motivated. This contrasts with proposals like FUSS according to which graded effects should be observed only for certain categories of knowledge. For actions, instead, the first aim of this experiment is to assess whether the basic semantic blocking effect in the object domain is also observed in the action domain. Also of interest is whether graded effects are observed for actions, and whether the degree of gradation differs for the two domains.

Method

Participants

Ninety-four native English speakers from the UCL community participated in the experiment in exchange for payments of £3. All had normal or corrected-to-normal vision.

Materials

Groups of action and object pictures were selected based upon FUSS semantic distance rather than on predefined categories such as tools, animals, etc. This was done beginning with all of the picturable words referring to objects and actions included in FUSS. Objects and actions were considered separately in this process as these semantic domains are largely separate (see chapter 5). From these sets of items, subsets were selected which exhibited both within-group semantic similarity (low semantic distances among them) and dissimilarity to other sets (high semantic distances between exemplars of different sets). In order to allow the investigation of graded similarity effects between sets of items, three sets of objects and three sets of actions (each containing eight items) were selected: two sets relatively close to each other, and a third relatively far from the other two.

For objects, the sets of items came from three well-defined categories with obvious category labels: vehicles (average within-set distance = 2.71 units), clothing (5.35), and body-parts (7.35). Clothing and body-parts were "near" sets (with an average distance between exemplars from the two = 13.51 units) while the other two were "far": vehicles and clothing (18.30) and vehicles and body-parts (18.53). Most object pictures were taken from Snodgrass & Vanderwart (1980) with a few prepared in a similar style specifically for this experiment.

For actions, the groups of pictures do not fall into such clearly-defined categories, but can be broadly designed as "body actions" such as *hop, kick, walk* (with an average within-set distance = 7.66), "tool actions" such as *cut, draw, shovel* (11.44), and "actions involving the mouth" such as *drink, frown, yawn* (12.21). Body and tool actions were "near" sets (16.74); body actions and actions involving the mouth (21.60), and tool actions and actions involving the mouth (20.45) were "far". Action pictures were taken from Druks and Masterson (2000), and additional pictures were drawn by the same artist who drew the pictures for Druks and Masterson. Semantic distances between the words referring to actions were somewhat larger than between the words referring to objects. This was necessary in order to ensure that the action pictures were distinguishable from each other. All items included in the experiment are listed in Appendix F.

<u>Visual similarity ratings</u>. Because semantic distance is correlated with visual similarity (see Vitkovitch, Humphreys & Lloyd-Jones, 1993), it was important to consider if pictures in these sets (which differ in semantic similarity to each other) also

differ visually. Visual similarity ratings were collected for the object and action pictures (following the procedure used by Damian et al., 2001), by presenting all possible pairs of object or action pictures to participants who were asked to rate their similarity in appearance. Instructions emphasised the focus upon visual appearance over semantics, providing examples such as *tennis racquet*, *guitar* and *piano* where visual similarity was high for members of different semantic categories, and low for members of the same semantic category. The scale ranged from 1 (not similar at all) to 5 (very similar); see Table 10 for average ratings as a function of semantic condition (within-set, e.g. visual similarity between pairs of body parts; near, e.g. visual similarity between body parts and clothing; far, e.g. visual similarity between body parts and vehicles). Most of the conditions differed significantly from each other; for objects, within-set pictures were the most visually similar to each other, followed by close pictures, and far pictures the least visually similar. For actions, within-set pictures were the most similar to each other, but close and far did not differ. However, visual similarity was rated as very low overall regardless of condition, reducing the likelihood that any putative semantic effect is due to visual similarity.

Table 10. Average visual similarity ratings between object and action pictures as a function of semantic condition (standard deviations in brackets).

	Semantic condition		
	Within-set	Close	Far
Objects	1.93 (1.14)	1.60 (0.97)	1.27 (0.87)
Actions	1.71 (1.20)	1.60 (1.04)	1.50 (0.96)

Preparation of experimental lists. Parallel experimental lists were prepared for objects and actions, including three experimental conditions created from combinations of the different sets of items (illustrated below with object-noun examples; action-verb lists were created in the same way): semantically far, semantically close, and within-set. Blocks of items representing the semantically far condition were created by selecting four items from each of two semantically far sets, e.g. {*arm, finger, foot, hand; aeroplane, bicycle, bus, car*}. Because there were eight items in each set, two such blocks were created (in this case, the second one would include {*leg, neck, shoulder, thumb; helicopter, lorry, motorcycle, train*}). The particular items that would appear in these two blocks were selected randomly for each participant. The semantically close condition was created in a similar manner, selecting four items from each of two semantically close sets (e.g. {*arm, finger, foot, hand; belt, glove, hat, shirt*} in the first set, and {*leg, neck, shoulder, thumb; shoe, sock, trousers, waistcoat*} in the other. Finally, the within-set condition was created using only members of a single set (e.g. {*arm, finger, foot, hand; leg, neck, shoulder, thumb*}). These conditions are summarised in Table 11:

Table 11: Composition of the different semantic conditions for words referring to objects and to actions, Experiment 5. The number of different blocks in a given condition appears in parentheses.

	Objects	Actions
Semantically far	Body parts and vehicles (2) Clothing and vehicles (2)	Body actions and mouth actions (2) Tool actions and mouth actions (2)
Semantically close	Body parts and clothing (2)	Body actions and tool actions (2)
Within-set	Body parts (1) Clothing (1) Vehicles (1)	Body actions (1) Tool actions (1) Mouth actions (1)

As can be seen from Table 11, there is only one instance of each within-set block, compared to the semantically close and semantically far conditions. In order to make these conditions statistically comparable to each other, two versions of each within-set block were created. In both cases, they contained all eight exemplars of a set, but half were treated as fillers: for example, among {*arm, finger, foot, hand; leg, neck, shoulder, thumb*}, in a first block {*arm, finger, foot, hand*} would be treated as experimental items (named in the context of other body parts) and as filler items in the second. The experimental items in the second block would be {*leg, neck, shoulder, thumb*} (again, named in the context of other body parts), and those items would be treated as fillers in the first block. This allows identical treatment of the three semantic conditions for each participant: four pictures from a set, named in the context of other items from within the set, in the context of items from a semantically near set, and in the context of items from a semantically far set. This resulted in 12 different blocks for object-nouns and 12 for action-verbs.

Within a block, each of the eight pictures was presented to be named a total of four times (32 trials per block). Pictures were ordered pseudorandomly: each picture was sampled once before any picture was repeated in a block, and no picture appeared twice in succession.

Each block was presented twice in the course of the experiment (thus a total of 24 blocks). Blocks were presented in pseudorandom order (each block was sampled once before any block was repeated in the experiment, and no block appeared twice in succession). Stimuli were presented using E-Prime experimental software (Schneider, Eschman & Zuccolotto, 2002) on IBM-PC compatible computers; response latencies were collected using a PST Serial Response Box (Psychology Software Tools) and tape recorded for error analysis.

Procedure

Participants were assigned to a word type condition (object-noun naming, n = 40; or action-verb naming, n = 54). They were instructed that pictures would be presented on the computer screen, and their task was to name the picture aloud as quickly as possible. They were asked to name the object pictures using single nouns and the action pictures, using the –ing form of the verb. Prior to the experiment proper, to ensure that they knew the target label of all the pictures, each picture was presented for the participants to name, and the experimenter provided the label in the (few) instances in which participants failed to name a picture correctly. Then, participants were presented with a practice block.

In the experiment proper, each block began with a button press of the participant. A fixation cross appeared on screen for 300ms, followed by a blank screen of 450ms. The target picture appeared in the centre of the screen and remained until the voice key detected a response, or 2500ms if no response was detected. Responses were followed by a 200ms blank screen, followed by the fixation cross for the next trial. Each session was tape-recorded and scored for accuracy.

Results

Response latencies

Error trials and response latencies faster than 250ms or slower than 1500ms were excluded from the latency analyses. Figure 6 reports average response latencies for object-nouns and action-verbs in the different semantic conditions.

Figure 6. Average correct naming latencies as a function of word type and semantic condition, Experiment 5.

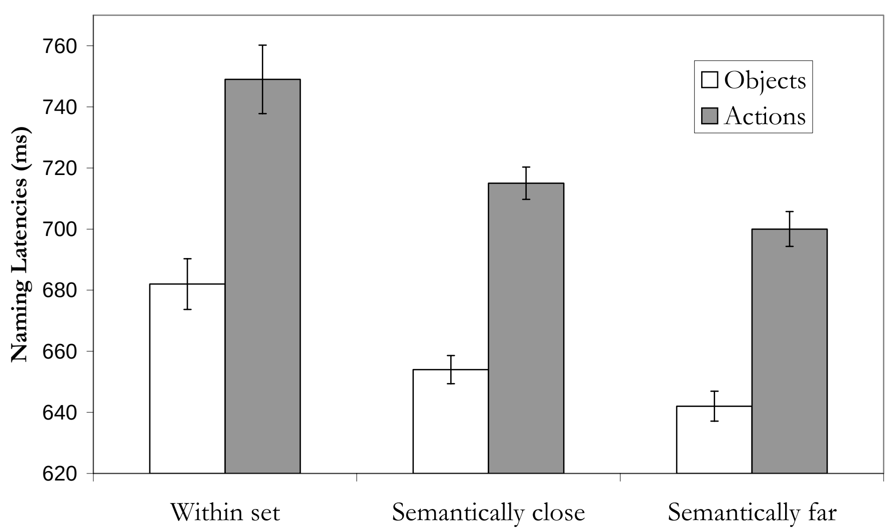

First, a two (word type: objects vs. actions) by three (semantic condition: within-set, semantically close and semantically far) ANOVA was carried out, both on subjects and items The main effect of word type was significant ($\underline{F}1(1,92) = 11.61$, $\underline{p} < .001$, $\underline{F}2(1,30) = 14.54$, $\underline{p}<.001$), reflecting longer naming latencies for actions than objects. The main effect of semantic condition was also significant ($\underline{F}1(2,184) = 19.01$, $\underline{p} < .001$, $\underline{F}2(2,60) = 16.14$, $\underline{p}<.001$), reflecting the effects of semantic distance. The interaction was not significant ($\underline{F}1(2,184) = 1.34$, $\underline{p} = .26$, $\underline{F}2(2,60) < 1$), indicating that the semantic condition had the same effect for object and action domains.

Since semantic groups were at different distances for objects and actions, separate linear trend analyses were performed on the simple main effects of semantic blocking to assess the role of semantic distance. Linear contrasts were calculated on the basis of the semantic distances between items in within-set, semantically close, and semantically far conditions, as intervals were not equidistant. The resulting contrast coefficients were [-6.4, 0.8, 5.6] for objects, and [-6.2, 1.0, 5.2] for actions. The linear trend was significant for objects ($F1(1,39)=20.87$, $p<.001$, $F2(1,15)=18.30$, $p<.001$); it was also significant for actions ($F1(1,53)=24.44$, $p<.001$, $F2(1,15)=21.73$, $p<.001$). Corresponding quadratic trends were not significant ($\underline{F}s < 1$). In both cases this pattern of data reflects a graded increase in naming latencies, with the fastest latencies observed in the semantically far condition, medium latencies in the semantically near condition, and the slowest latencies in the within-set condition.

<u>Errors</u>

Errors occurred on 6.9% of the trials and involved failure to detect initial word onset, cases in which the voice relay detected sounds before the initial word onset, and erroneous or dysfluent utterances. Analysis of variance showed no significant effect of semantic blocking condition on the number of errors, either for object or action naming (all $\underline{F}s < 1$).

Discussion

Parallel, graded patterns of semantic blocking interference were observed for separate sets of words referring to objects and to actions, selected not on the basis of *a priori* category membership, but from FUSS semantic distances. These results cannot be attributed solely to visual similarity, which is correlated to semantic similarity, mainly because graded semantic effects were observed for action-verbs even though visual similarity ratings did not significantly differ for semantically near and semantically far conditions. If visual similarity were playing a role in these results, it might be expected that word type would interact with semantic condition (because object-noun pictures were rated as more visually similar for semantically near than for semantically far items).

These results are important because they complement the results presented in Chapter 6 which showed that FUSS can predict graded semantic effects between pairs of highly-related words. FUSS can also predict graded semantic effects between groups of words; thus its representations also accurately reflect a higher level of organisation, in addition to being indicative of similarity among individual words. Most important is the fact that these effects are not only observed for words referring to objects for which category-level organisation is quite clear, but also for words referring to actions where this is less so. These results also are contrary to theories of organisation that do not permit gradation between categories (e.g., Caramazza & Shelton, 1998), but instead are consistent with views in which semantic fields in a given domain (object, action) are not strictly distinct from each other (consistent with the notion that category boundaries are only vaguely defined, perhaps best described in terms of family resemblance; Rosch & Mervis, 1975).

Chapter 9: Discussion and conclusions

General summary

Previous chapters describe the implementation and testing of FUSS (Featural and Unitary Semantic Spaces), made up of a conceptual representational space which is operationalised by the speaker-generated features themselves, and a separate level of semantic representation, operationalised as a space derived from properties of similarity among words' featural content. Crucially, this model included representations not only of nouns referring to objects, a domain which has already received substantial attention using similar approaches, but also nouns and verbs referring to actions, using a common set of representational assumptions across domains.

This work began with the collection of speaker-generated features for a collection of words referring to objects and actions, described in Chapter 2. Analysis of featural content (Chapter 3) reveals how the broad domains of objects and actions differ from each other, but also highlighted substantial differences between semantic fields within each domain. The distribution and patterns of feature types across semantic fields are consistent with a wide range of evidence from a number of research domains. For example, feature type composition of living things vs. artefacts converge with results of imaging studies showing that sensory-motor areas are differentially activated for these types of words (e.g. Martin & Chao, 2001; Hauk et al., 2004; Tettamanti et al., 2005; Vigliocco et al., 2006). They can also account for patterns of category-related impairment after brain damage (particularly the living/nonliving distinction) in terms of feature types, also making additional novel predictions about the extent to which other domains of knowledge should also be affected. For example, words referring to actions related to sensory experience (e.g. light emission, sounds and noises) should be impaired along with living things if such impairments hinge on dependence on sensory features.

As a next step, the speaker-generated features were used to generate a separate level of lexical-semantic representation, modelled using self-organising maps. Properties of similarity between words' featural input led to patterns of organisation at this level as described in Chapter 4, ranging from the broad distinction between objects and actions, to category-level organisation, to fine-grained organisation among individual words in a semantic field (Chapter 5). Differences between words referring to objects and actions, often used as motivation for entirely separate representation schemes, emerged despite a single set of assumptions underlying the representation of both in FUSS. For example, words referring to objects exhibited strong categorical distinctions for the most part, while this was seldom true for words referring to actions, a difference emergent from characteristics of the featural input rather than any differences in the representational space *per se*. These analyses also revealed that words' grammatical class did not exert any effects on their semantic representations beyond the semantic distinction between objects and actions: nouns referring to actions were not distinguished at the lexical-semantic level from verbs with similar meanings.

In a series of experiments, behavioural consequences of FUSS's semantic similarity were tested, using tasks where semantic effects arise automatically in language processing. FUSS measures of fine-grained similarity among pairs of words predicted the degree of semantic priming in lexical decision, both for objects (Chapter 6, Experiment 1) and for actions (Chapter 6, Experiment 2), and predicted the degree of interference from distracter words in picture naming, again both for objects (Chapter 6, Experiment 3) and actions (Chapter 6, Experiment 4). Crucially, comparison of FUSS with other models of semantic representation from which word-level similarity measures can be derived (LSA, Wordnet) revealed that FUSS outperformed the other two models across these four data sets (Chapter 7). Predictions derived from FUSS also predicted effects beyond the word level: graded semantic blocking effects in picture naming were observed based on FUSS relatedness of sets of words, again both for objects and actions (Chapter 8, Experiment 5). Together these results show how the assumptions of FUSS lead to properties of lexical representation and processing that are consistent with a wide range of data from behavioural experimentation, patterns of impaired performance after brain damage, and imaging results.

FUSS as a theory of lexical-semantic representation

We now return to the central questions raised in the Introduction. What is the content of word meaning? How is word meaning organised? What is the relation between words? In FUSS, it is clear what makes up word meaning: first, featural properties of meaning corresponding to concepts, some of which are organised according to the sensory/motor channel by which they are experienced (e.g. visual, auditory, tactile, motoric, etc.). This conceptual level of representation is thus organised by modality, although any given concept should be considered as a group of coactivated features across multiple modalities. Speaker-generated features serve as a verbal proxy for this input and allow us to investigate how different words' meanings depend upon different input channels. Words' meanings are represented separately, in a supramodal system that serves to integrate information across modalities, a system that further serves to bridge between concepts and lexical information such as syntax, orthography and phonology. This system is organised according to properties of similarity among sets of features for different words. Implemented here using self-organising maps (Kohonen, 1997), similarity among words emerges on the basis of characteristics of the featural input such as shared and distinctive features, feature salience and patterns of correlation and decorrelation among features, expressed in terms of spatial proximity in a low-dimensionality space ("map") derived from the high-dimensionality featural input.

Of course, as described here, there are several important potential weaknesses of FUSS which should be addressed at this stage. Foremost among these is the verbal nature of the featural input: all of the input to FUSS as implemented here is speaker-generated features: written words or phrases referring to properties that make up meaning. After all, the actual input to the system is meant to be via sensory-motor systems rather than verbal descriptions. As such it might be argued that these features merely reflect verbal knowledge (thus, information represented at the lexical-semantic level) rather than providing any insight into nonlinguistic (i.e. perceptual/motoric) information. This can be addressed, however, by reference to the imaging literature; there are now a large number of studies showing that processing words produces activation in sensory-motor areas that correspond very well to the featural makeup of such words in FUSS's feature set (e.g. Martin & Chao, 2001; Hauk et al., 2004; Tettamanti et al., 2005; Vigliocco et al., 2006). If the sensory-motor properties of speaker-generated features were only informative about verbal knowledge and not about sensory-motor knowledge, such correspondence would not be expected. Featural makeup in FUSS also corresponds well with patterns of impairment that have often been attributed to differential impairment to sensory or functional properties of objects (particularly the distinction between living and nonliving entities), e.g. (SFT cites).

There is also, however, an additional concern also related to the verbal nature of the speaker-generated features: there are certain kinds of featural input difficult to describe verbally (at least in English). This was particularly noticeable when participants attempted to describe shapes and sizes of objects. Beyond simple descriptors of shape (e.g. <round>, <straight>, <long>, <thin>) and size (e.g. <big>, <small>), participants diverged greatly in the way they attempted to express finer details of shape and size, nearly all of which had to be discarded as "idiosyncratic" given the feature collection methods, despite containing important information clearly relevant to the objects in question. For example, some participants attempted to estimate size precisely (e.g. *hammer* is <between six and 12 inches long, and about one inch wide>, *apple* is <about six inches around>), others relatively using arbitrary anchors (e.g. *cat* is <smaller than a dog> and <larger than a mouse>, *elephant* is <larger than a person>), and others in terms of relative dimensions (e.g. *hammer* is <wider than it is long>). Concerning shape, a similar range of strategies were observed. Some participants attempted to describe objects using idiosyncratic descriptions of combinations of simple verbalisable shapes (e.g. *scissors* has <two rings, attached to two long triangles, joined together part way down>, *pear* is <like a circle with a bump on top>), others used shape analogies (e.g. *drill* is <shaped like a gun>), and a number of participants reported after the experiment that they would have liked to describe the shapes of some things but were unable to do so clearly (e.g. a few instances of features like <shape?> which appeared to reflect the same kind of difficulty). Issues like these could result in problems for representations of those words for which properties like size and shape are important, under-representing them due to difficulty in verbalising them. However, this potential problem also seems to have limited consequences. This can be seen from the results of the behavioural experiments presented in Chapters 6 and 8. If some types of words have inaccurate or limited featural representations, this should have translated into poor performance by FUSS in predicting fine-grained behavioural effects. Instead, FUSS consistently exhibited strong performance in predicting the results of multiple tasks involving words referring to objects and to actions, and did not appear to be particularly worse for any of the semantic fields from which words were selected for these experiments. Presumably, then, other properties that participants were able to verbalise were sufficient to make up for this particular limitation of FUSS.

Another potential criticism of FUSS applies to its apparent inflexibility in representation. After all, the featural input corresponding to each word is a fixed vector of feature weights, and each word's representation at the lexical-semantic level corresponds to a single point in representation space (defined via ensemble averages across multiple self-organising maps). Words, however, are used flexibly, as illustrated for example by the existence of strong context effects. To exemplifiy the importance of context effects, Barclay, Bransford, Franks, McCarrell, and Nitsch (1974), in a classic experiment showed that when participants saw words like *piano* in sentences that either stressed a piano's weight ("The man lifted the piano") or its sound ("The man tuned the piano") their recall of the word *piano* was better in response to a cue that was related to the original context: *heavy* cued *piano* better than *with a nice sound* when the weight had been stressed in the initial sentence; and the reverse pattern was true when the sound had been stressed. This pattern of results is difficult to accommodate within a binary featural system. The differential effects of context seem to require that features like <heavy> and <makes sound/music> can have differential weights, depending on the context. However, speaker-generated features used in FUSS were collected in a neutral context, in lists consisting of other (unrelated) words, suggesting that FUSS may only be capable of representing words in neutral contexts (or prototypical contexts). Properties of FUSS, however, enable it to deal straightforwardly with flexibility in representation and context effects. First, although this particular implementation of FUSS was trained with a single training vector with fixed weights for each word in the set, a more realistic model of language experience would instead include a variable set of featural inputs corresponding to the features that are important for that word in that particular episode. This serves to set up distributional conditions making it more likely that a new instance will be encompassed or at least near existing regions of conceptual/feature space that map onto words in lexical-semantic space. For example, exposure to *piano* in the two different contexts above would include many features in common, but very different weights related to its weight or its sound. Over multiple exposures to different sets of featural inputs, a region of multidimensional featural space will come to be associated most strongly with *piano* and thus any subsequent activation falling within this region of feature space will also correspond to *piano*. This does not even require that the model has received input spanning the entire potential range of contexts, but even permits it to respond correctly in novel contexts. This is because the algorithms underlying self-organising maps are designed to select a "winner": the best-matching output unit (here, "word"), corresponding to any possible input in feature space. In other words, effects of context such as those of Barclay et al. (1974) would arise as variation from the prototypical weights of a word's features, which nonetheless are still most similar to that word than to any others, according to the model's representation state at that point in time.

One other potential concern reflects FUSS's ability to represent other domains of knowledge. Although FUSS represents an important move forward from models that concern only words referring to concrete objects, it still has a highly limited range of domains (some objects, actions and events). Although the success of FUSS at representing action words as well as object-nouns should not be underestimated, future work should also consider other domains of words. For example, consider the case of properties and qualities (somewhat corresponding to adjectives and adverbs) – data from speech errors (Garrett) and semantic field analysis suggest these are typically organised around poles of opposition which could be roughly considered to be something like basic level for objects, e.g. *large-small*; *smooth-rough*; *dark-light*; *fast-slow*. Some properties (of various kinds) were included in a pilot stage of the feature collection phase, and participants tended to produce features of the following sorts. First, superordinate features were quite common, such as <size>, <texture>, <brightness>, <speed>; also common were features referring to the sensory channel through which a property could be experienced, such as <vision>, <touch>,

<hearing>. Both features referring to superordinates and sensory channels are also very typical of object and action words. For properties and qualities, participants were also very likely to generate lists of entities for which a given property is prototypical (e.g., perhaps <elephant>, <house>, <tree> for *large*; <sun>, <light bulb>, <fire> for *bright*), and they were also much more likely to produce antonymic features (e.g. <not large> for *small*). Within certain relational accounts such as network models, both of the latter would straightforwardly correspond to types of labelled links between concepts. In FUSS, instead, patterns of lexical-semantic similarity would result from sets of properties that are prototypical for the same entities, in addition to shared superordinate and sensorimotor features. Of course this cannot be the end of the story, because mutual relationships of various kinds are lost. For example, the feature <not black> is not linked to the word *black*, nor is there a mutual relationship between the feature <red> of the word *apple*, and the word *red* with feature <apple>, or vice versa (see also Hampton, 1981, for a discussion of how more complex relationships between features may be needed when considering featural representations of abstract concepts). Nonetheless, other shared properties should be sufficient to provide a high quality estimate of semantic similarity among words of these kinds. This suggests that FUSS can be a promising approach beyond objects, actions and events.

FUSS in the context of other theoretical approaches

Concerning the various theoretical approaches to representing meaning, FUSS can be seen as a sort of hybrid, incorporating both elements of featural views and elements of relational views. This dual nature is in contrast to many models of meaning, as it arises because FUSS draws a strong distinction between conceptual and lexical-semantic levels of representation, in each of which different principles are instantiated. Featural views are reflected straightforwardly in the model's reliance upon speaker-generated features to reflect the conceptual level of representation (e.g. Rosch & Mervis, 1975; Smith, Shoben & Rips, 1974; Collins & Quillian, 1969; Hampton, 1979; 1981; Hampton & Gardiner, 1983; Jackendoff, 1990; Minsky, 1975; Norman & Rumelhart, 1975; Shallice, 1993; Smith & Medin, 1981; and on larger scales by Cree & McRae, 2003; Cree et al., 2006; Cree et al., 1999; Garrard et al., 2001; McRae & Cree, 2002; McRae et al., 2005; McRae et al., 1999; Randall et al., 2004; Rogers & McClelland, 2004). Like many of the featural models developed in the 1970s, a strength of FUSS derives from its flexibility. Word meaning is not considered to be strictly based upon certain features, but instead upon sets of co-occurring features, weighted according to salience, and considering a word to correspond to a probabilistic volume in multidimensional feature space rather than a single point (akin to views permitting fuzzy category boundaries, and using principles of family resemblance, (e.g. Hampton, 1979; Rosch, 1973; Rosch & Mervis, 1975; Rosch et al., 1976). FUSS goes beyond typical featural models developed in the cognitive tradition, however, by explicitly linking some of these features to sensorimotor experience, and as such grounding representation in reality. As such it is strongly influenced by models developed in cognitive neuroscience and neuropsychology (e.g. Farah & McClelland, 1991; Devlin et al., 1998)..

But FUSS is also relational in nature, when it comes to the lexical-semantic level. Here, semantic similarity measures are strictly relational in nature, corresponding as they do to ensemble average distances between points in the multiple self-organising maps that reflect this level of representation as implemented here. Crucially, and unlike many relational models (e.g., HAL: Burgess & Lund, 1997; LSA: Landauer & Dumais, 1997; Osgood, 1962; Osgood et al., 1975; Osgood et al., 1957; Snider & Osgood, 1969), these similarity measures are interpretable, corresponding as they do to aspects of the featural input. This also obviates a major concern that has generally been levelled at relational models: that they do not tend to be grounded in reality in any way. Here, the relational lexical-semantic level is grounded in reality through its featural input. The particular relations between items in FUSS's lexical-semantic space, however, cannot be directly analysed in the same manner as is possible with relational models where the relations are contentful, such as network models where relational links are labelled (e.g. Collins & Loftus, 1975; or Wordnet: Miller & Fellbaum, 1991), or semantic field theory where relations are straightforwardly described in terms of specific principles that are relevant only to certain fields of knowledge.

Grounding language in experience

One of the fundamental assumptions underlying FUSS is the importance of sensorimotor experience to the development and representation of word meaning. In this particular implementation, such experience is represented by verbal speaker-generated features referring to information gained through sensorimotor channels. Although it is uncontroversial that word meaning is learned to an important extent through interactions with the world through sensorimotor experience, various theories differ with respect to extent to which linguistic representations and processes are linked with sensory and motor representations and processes. As discussed by Meteyard & Vigliocco (in press), classes of theories of semantic representation can be considered to fall along a continuum in this regard. At one end of the continuum are strong embodiment hypotheses which embed assumptions of necessary and direct engagement: semantic representations of (concrete) entities and events necessarily depend upon primary sensory and motor systems, and that those systems are directly engaged during semantic processing rather than being transduced or mediated by other systems. The tight link and shared characteristics of the two systems would create strong dependency relations between the two: semantic processing would necessarily engage sensorimotor systems, and vice versa. Such theories are exemplified by Gallese and Lakoff (2005) who propose that all aspects of semantic representation and processing are contained across multimodal sensorimotor systems rather than being reduced to a common (amodal or supramodal) system. Similar in character are theories by Pulvermüller (2001), Barsalou and colleagues (Barsalou, 1999; Barsalou, Kyle Simmons, Barbey & Wilson, 2003), and Glenberg and colleagues (Glenberg & Robertson, 2000; Glenberg & Kaschak, 2002, 2003) all of whom propose that semantic representations and processes are automatically and necessarily linked to low-level sensory and motor systems (see Meteyard, 2008; Meteyard & Vigliocco, in press). At the other extreme of this continuum are purely symbolic, amodal theories where semantic representations are fully independent from sensorimotor content, with any links between the two occurring outside the semantic system. Any exchange of information between the two systems would thus be necessarily mediated by other cognitive systems. Examples of such theories include the WEAVER++ model of lexical retrieval (Levelt, 1989; Levelt et al, 1999), global co-occurrence models such as LSA (Landauer & Dumais, 1997) and HAL (Burgess & Lund, 1997), and relational similarity accounts as described by Osgood and colleagues (Osgood, 1962; Osgood et al., 1975; Osgood et al., 1957; Snider & Osgood, 1969).

The two contrasting classes of theories above, however, are not the only possibilities, but merely the extremes of a continuum. For example, "weak embodiment" hypotheses (Meteyard, 2008; Meteyard & Vigliocco, in press) are those in which semantic representations are grounded in low-level sensorimotor systems, but the semantic system itself is supramodal, serving to bind together information from different modal systems. FUSS falls into this class of theories: semantic representations (whether implemented as self-organising maps as described in Chapter 4, or any other implementation whereby this level of representation is conceived as a single representation space uniting featural information across modalities) are supramodal but grounded in low level sensorimotor systems. Like the strong embodiment hypotheses, these theories hold that sensorimotor systems are essential underpinnings of semantic representations, but they diverge in that these systems are not so strictly linked, either in terms of the necessity for these systems to be invoked in all semantic processing, or in terms of the direct processing links between the two (the degree of directness thus defining a particular theory's position along the continuum between strongly embodied and abstract/amodal). Effects showing interdependence between semantic and lower-level sensorimotor systems should be observed under some conditions (depending upon the processing assumptions of a particular model), but unlike strong embodiment hypotheses such effects need not necessarily be symmetrical nor arise in all circumstances.

There is now a substantial and growing body of evidence concerning the relation between low-level sensorimotor systems and language processing, generally favouring embodied hypotheses over strictly amodal ones (see Meteyard & Vigliocco, in press, for an extensive review). In behavioural experimentation, evidence of this nature can be seen in studies showing that "purely linguistic" tasks are affected by perceptual or motor processing, or that "purely perceptual" or "purely motor" tasks are affected by semantic content of language stimuli. Such findings are inconsistent with the central assumptions of amodal theories, under which these kinds of processes should be independent of each other, but fall naturally from embodied assumptions under which they are closely linked. For example, a close link between motor systems and sentence processing systems has been demonstrated in a number of studies. For example, Glenberg and Kaschak (2002) presented participants with sentences describing motor actions in the imperative form (e.g. "Close the drawer"), or as transfer actions involving themselves (e.g. "Courtney handed you the notebook") and asked them to judge their sensibility by button presses. Crucially, the buttons were laid out in a configuration that required participants to move their hand either away from or toward their body.

Participants were faster to respond when the direction of response was consistent with the physical motion implied in the sentence content (e.g. moving the hand away from the body for "Close the drawer", or toward the body for "Courtney handed you the notebook") than when it was inconsistent (see also Borreggine & Kaschak, 2006). Such results should not be observed given the assumptions of amodal systems, because the type of response should be irrelevant to the task of deciding whether a sentence makes sense or not if semantic processing is independent from motor systems. Many other studies have shown the same kinds of effects using a range of motor responses: responding with hand vs. foot in judging sentences involving hands or feet (Buccino, Riggio, Melli, Binkofski, Gallese, & Rizzolatti, 2005); responding with rotary motion to sentences implying rotation in a particular direction (Zwaan & Taylor, 2006); and performing manual sorting tasks involving directional motion while producing sentences referring to directional motion (Casasanto & Lozano, 2007).

Results consistent with embodiment have also come from tasks related to visual perception of motion. For example, Meteyard, Bahrami and Vigliocco (2007) asked participants to do a difficult perceptual task involving identification of coherent visual motion near threshold, and at the same time they were listening to blocks of words referring to directional motion not relevant to the task. When the two were inconsistent (e.g., visual motion was upwards and words referred to downwards motion), participants were less able to detect coherent visual motion, reflected in reduced d'. This indicates effects of (passive) lexical processing on a low-level perceptual task, consistent with embodiment hypotheses where these would be closely linked. Meteyard, Zokaei, Bahrami and Vigliocco (in press) also found effects of motion perception on language processing: threshold-level patterns of directional motion impeded lexical decision on words referring to direction when the two were inconsistent. Any number of findings consistent with embodiment hypotheses can also be found in the imaging literature (reviewed in Meteyard & Vigliocco), where it has long been taken for granted that meaning must be grounded in sensorimotor experience, in part because neuroscientific theorising mostly does not make a distinction between conceptual and semantic representations.

Together these findings provide strong evidence against strictly amodal accounts of semantic representation, but are not definitive with respect to contrasting strong from weak embodiment hypotheses. One possible angle for gaining leverage on this issue may come from the distinction between semantic and conceptual levels of representation. As discussed in the introduction, cross-linguistic evidence seems to demand not only that these levels of representation be distinguished from each other, but also that they embed (at least some) different principles of organisation. This is particularly evident from the results of Vigliocco, Vinson, Paganelli and Dworzynski (2005) and Kousta et al (in press), where effects of Italian gender are observed at the semantic level but not at the conceptual level. Under weak embodiment hypotheses like FUSS, such findings can be easily accommodated due to the different principles of organisation of the (modal) conceptual system and the (supramodal) lexical-semantic system. For strong embodiment theories, however, these findings seem to demand that some differences exist between semantic and conceptual systems, an arrangement difficult to implement given the strong dependency between these both in terms of representation and processing.

Going beyond purely sensory and motor information

Up to this point, the discussion of FUSS has revolved around the extent to which its representations can be derived from sensorimotor input, without much regard for the other kinds of information that make up a sizeable fraction of the speaker-generated features that are produced across different types of words. As discussed in

Chapter 3, "Other" features, those which did not fall into any category of sensory, functional or motoric, account for 37.6% of all feature weights among the words in the set, and were even higher in some domains of knowledge (e.g. half of all feature weights for words referring to clothing and communication were classified as "Other"). Such features reflect a wide range of types, including encyclopaedic information (e.g. that zebras and elephants come from Africa; that tomatoes grow in gardens and onions grow in the ground; that cows give milk and live on farms); compositional information (e.g. that combs are made of plastic, and daggers are made of metal); information about the kinds of participants that can perform a particular action (e.g. that punching and arriving are done by humans, while licking, drinking and tasting are done by humans and by animals); information about superordinate category labels (e.g. that dogs are pets, mammals and animals; or that knives are utensils, tools, weapons and objects); information about higher-order cognitive processes (e.g. that giving is intentional and involves generosity; or that hiccuping is involuntary, embarrassing and disruptive), and many other sorts of features that also do not easily fall into sensorimotor classes. In fact, a large proportion of "Other" features seem to be best described as being learnt through experience with language much more than direct sensorimotor experience. It is therefore important to consider how such information could come to play a role in semantic representation, particularly given the crucial role of sensorimotor experience as discussed above, and the apparent need for a linguistic system to be in place before language information can contribute meaningfully.

A developmental framework within which this apparent paradox can be explained has been advanced by Gleitman and colleagues (Gleitman, 1990; Gleitman, Cassidy, Nappa, Papafragou & Trueswell, 2005; Landau & Gleitman, 1985; Trueswell & Gleitman, 2004). Under this account, multiple sources of information contribute to the development of the lexicon. Different sources contribute differentially depending on the domain of knowledge, and the availability of these different sources evolves over time. More specifically, the initial state is one where the only available information is the extralinguistic context accompanying a word, sensorimotor information in other words. This serves well for the sensorimotor properties of concrete objects for which such properties are especially salient, but less so for actions and even less so for more abstract concepts ("hard words", Gleitman et al. 2005). Once these basic-level object representations develop, they serve as a foundation for the development of syntactic knowledge, which in turn can contribute to further lexical development (e.g., learning action verbs). More abstract words would be learnt through a similar developmental process, as metalinguistic knowledge builds even further upon ordered, contentful linguistic content. In FUSS, representing a mature (adult) system, this developmental trajectory is no longer evident. Instead, all of these multiple sources of information are in place and contribute to word meaning appropriately depending on the domain of knowledge.

Current research in our lab is exploring these issues in more depth, assessing the extent to which the development of semantic representations can benefit from interaction between sensorimotor information and linguistic contexts, rather than a system which treats the two as independent sources of information (Andrews, Vigliocco & Vinson, 2005a, b; submitted). We are also beginning to investigate the meanings of abstract words, which are not included in the implementation of FUSS presented here. After all, any theory of word meaning should also account for representations of abstract words, which have been relatively neglected in theories of semantic representation. It is often considered that abstract words are solely or mostly represented linguistically while concrete words are also grounded in perception and action. Although this is most evident in dual coding theory (Paivio, 1971; 1986; 1991; 2007) where concrete words benefit from access to a nonlinguistic "imagistic" system, such views are widely prevalent (see Kousta, Vinson, Andrews & Vigliocco, submitted, for a review). However, initial evidence is strongly suggestive that, just like concrete words, the representations of abstract words may also develop through links with lower-level systems—those involved in processing emotions (Barsalou & Wiemer-Hastings, 2005; Kousta et al., submitted). These lines of work should further illuminate the extent to which sensorimotor, emotional and linguistic input converge in providing input to developing the meanings of words.

Appendix A. Items for which speaker-generated features were obtained, and their semantic field labels. Words referring to actions are classified generally following Levin (1993).

Complete featural data, including feature weights for each word, and feature type classification as described in Chapter 3, have been permanently archived online, and may be downloaded from www.psychonomic.org/archive (search for Author: Vinson; Year: 2007).

the loss	noun: action
the ache	noun: body-action
the blink	noun: body-action
the burp	noun: body-action
the cough	noun: body-action
the frown	noun: body-action
the hiccup	noun: body-action
the itch	noun: body-action
the smile	noun: body-action
the sneeze	noun: body-action
the snore	noun: body-action
the squint	noun: body-action
the throw	noun: body-action
the touch	noun: body-action
the tremble	noun: body-action
the wink	noun: body-action
the yawn	noun: body-action
the pull	noun: change-location
the push	noun: change-location
the call	noun: communication
the challenge	noun: communication
the chat	noun: communication
the command	noun: communication
the cry	noun: communication
the demand	noun: communication
the plea	noun: communication
the request	noun: communication
the scream	noun: communication
the shout	noun: communication
the sigh	noun: communication
the suggestion	noun: communication
the threat	noun: communication
the whine	noun: communication
the whisper	noun: communication
the yell	noun: communication
the construction	noun: construction
the repair	noun: construction
the crash	noun: contact
the hit	noun: contact
the knock	noun: contact
the slap	noun: contact
the bombardment	noun: destruction
the destruction	noun: destruction
the murder	noun: destruction
the donation	noun: exchange
the exchange	noun: exchange
the loan	noun: exchange
the trade	noun: exchange
the flame	noun: light-emission
the flash	noun: light-emission
the flicker	noun: light-emission
the glow	noun: light-emission
the shine	noun: light-emission
the sparkle	noun: light-emission
the approach	noun: motion-direction
the arrival	noun: motion-direction
the ascent	noun: motion-direction
the descent	noun: motion-direction
the entry	noun: motion-direction
the escape	noun: motion-direction
the return	noun: motion-direction
the chime	noun: noise
the clang	noun: noise
the clash	noun: noise
the clatter	noun: noise
the crackle	noun: noise
the screech	noun: noise
the chirp	noun: noise-animal
the growl	noun: noise-animal
the meow	noun: noise-animal
the oink	noun: noise-animal
the bear	noun: animal
the bird	noun: animal
the camel	noun: animal
the cat	noun: animal
the cow	noun: animal
the dog	noun: animal
the donkey	noun: animal
the duck	noun: animal
the elephant	noun: animal
the fish	noun: animal
the fox	noun: animal
the giraffe	noun: animal
the goat	noun: animal
the horse	noun: animal
the leopard	noun: animal
the lion	noun: animal
the mouse	noun: animal
the pig	noun: animal
the rabbit	noun: animal
the sheep	noun: animal
the swan	noun: animal
the tiger	noun: animal
the wolf	noun: animal
the zebra	noun: animal
the ankle	noun: body part
the arm	noun: body part
the beak	noun: body part
the chin	noun: body part
the ear	noun: body part
the elbow	noun: body part
the eye	noun: body part
the face	noun: body part
the feather	noun: body part
the finger	noun: body part
the fur	noun: body part
the hair	noun: body part
the hand	noun: body part
the head	noun: body part
the knee	noun: body part
the leg	noun: body part
the lips	noun: body part
the mouth	noun: body part
the neck	noun: body part
the nose	noun: body part
the paw	noun: body part
the shoulder	noun: body part
the tail	noun: body part
the teeth	noun: body part
the thumb	noun: body part
the toe	noun: body part
the tongue	noun: body part
the wing	noun: body part
the wrist	noun: body part

the belt	noun: clothing
the blouse	noun: clothing
the coat	noun: clothing
the dress	noun: clothing
the glove	noun: clothing
the hat	noun: clothing
the mitten	noun: clothing
the pants	noun: clothing
the scarf	noun: clothing
the shirt	noun: clothing
the shoe	noun: clothing
the skirt	noun: clothing
the sock	noun: clothing
the suit	noun: clothing
the sweater	noun: clothing
the vest	noun: clothing
the apple	noun: fruit and vegetable
the artichoke	noun: fruit and vegetable
the asparagus	noun: fruit and vegetable
the banana	noun: fruit and vegetable
the bean	noun: fruit and vegetable
the broccoli	noun: fruit and vegetable
the cabbage	noun: fruit and vegetable
the carrot	noun: fruit and vegetable
the cauliflower	noun: fruit and vegetable
the celery	noun: fruit and vegetable
the cherry	noun: fruit and vegetable
the corn	noun: fruit and vegetable
the cucumber	noun: fruit and vegetable
the eggplant	noun: fruit and vegetable
the grape	noun: fruit and vegetable
the grapefruit	noun: fruit and vegetable
the lemon	noun: fruit and vegetable
the lettuce	noun: fruit and vegetable
the lime	noun: fruit and vegetable
the mushroom	noun: fruit and vegetable
the onion	noun: fruit and vegetable
the orange	noun: fruit and vegetable
the pea	noun: fruit and vegetable
the peach	noun: fruit and vegetable
the pear	noun: fruit and vegetable
the pepper	noun: fruit and vegetable
the pineapple	noun: fruit and vegetable
the plum	noun: fruit and vegetable
the potato	noun: fruit and vegetable
the pumpkin	noun: fruit and vegetable
the raisin	noun: fruit and vegetable
the raspberry	noun: fruit and vegetable
the spinach	noun: fruit and vegetable
the strawberry	noun: fruit and vegetable
the watermelon	noun: fruit and vegetable
the bomb	noun: misc. artefact
the book	noun: misc. artefact
the box	noun: misc. artefact
the carpet	noun: misc. artefact
the ceiling	noun: misc. artefact
the chair	noun: misc. artefact
the couch	noun: misc. artefact
the curtain	noun: misc. artefact
the door	noun: misc. artefact
the doorknob	noun: misc. artefact
the fence	noun: misc. artefact
the floor	noun: misc. artefact
the fork	noun: misc. artefact
the gate	noun: misc. artefact
the roof	noun: misc. artefact
the rug	noun: misc. artefact
the seat	noun: misc. artefact
the sofa	noun: misc. artefact
the stool	noun: misc. artefact
the table	noun: misc. artefact
the wall	noun: misc. artefact
the window	noun: misc. artefact
the axe	noun: tool
the broom	noun: tool
the brush	noun: tool
the chisel	noun: tool
the comb	noun: tool
the crowbar	noun: tool
the dagger	noun: tool
the drill	noun: tool
the dustpan	noun: tool
the file	noun: tool
the gun	noun: tool
the hammer	noun: tool
the hatchet	noun: tool
the hoe	noun: tool
the knife	noun: tool
the pen	noun: tool
the pencil	noun: tool
the pliers	noun: tool
the rake	noun: tool
the razor	noun: tool
the saw	noun: tool
the scissors	noun: tool
the screwdriver	noun: tool
the shield	noun: tool
the shovel	noun: tool
the spoon	noun: tool
the sword	noun: tool
the toothbrush	noun: tool
the tweezers	noun: tool
the wrench	noun: tool
the airplane	noun: vehicle
the bicycle	noun: vehicle
the boat	noun: vehicle
the bus	noun: vehicle
the car	noun: vehicle
the helicopter	noun: vehicle
the motorcycle	noun: vehicle
the raft	noun: vehicle
the ship	noun: vehicle
the train	noun: vehicle
the tricycle	noun: vehicle
the truck	noun: vehicle
the van	noun: vehicle
to find	verb: action
to lose	verb: action
to bleed	verb: body-action
to blink	verb: body-action
to breathe	verb: body-action
to burp	verb: body-action
to cough	verb: body-action
to cry	verb: body-action
to drink	verb: body-action
to drool	verb: body-action
to eat	verb: body-action
to feel	verb: body-sense
to frown	verb: body-action
to grin	verb: body-action
to hear	verb: body-sense
to hiccup	verb: body-action
to hold	verb: body-action
to inhale	verb: body-action
to inject	verb: body-action
to itch	verb: body-action

to kick	verb: body-action
to knock	verb: body-action
to lick	verb: body-action
to listen	verb: body-sense
to look	verb: body-sense
to notice	verb: body-sense
to retch	verb: body-action
to see	verb: body-sense
to sense	verb: body-sense
to shave	verb: body-action
to sit	verb: body-action
to smell	verb: body-sense
to smile	verb: body-action
to smoke	verb: body-action
to sneeze	verb: body-action
to sniff	verb: body-sense
to snore	verb: body-action
to spit	verb: body-action
to squint	verb: body-action
to stand	verb: body-action
to stay	verb: body-action
to swallow	verb: body-action
to taste	verb: body-sense
to throw	verb: body-action
to tickle	verb: body-action
to touch	verb: body-sense
to tremble	verb: body-action
to vomit	verb: body-action
to wash	verb: body-action
to watch	verb: body-action
to wink	verb: body-action
to yawn	verb: body-action
to ache	verb: body-sense
to die	verb: body-action
to carry	verb: change-location
to drag	verb: change-location
to drop	verb: change-location
to eject	verb: change-location
to lift	verb: change-location
to move	verb: change-location
to place	verb: change-location
to pull	verb: change-location
to push	verb: change-location
to put	verb: change-location
to send	verb: change-location
to bend	verb: change-state
to blend	verb: change-state
to empty	verb: change-state
to fill	verb: change-state
to mix	verb: change-state
to pour	verb: change-state
to shake	verb: change-state
to spray	verb: change-state
to stir	verb: change-state
to twist	verb: change-state
to admit	verb: communication
to advise	verb: communication
to argue	verb: communication
to ask	verb: communication
to call	verb: communication
to chat	verb: communication
to command	verb: communication
to demand	verb: communication
to greet	verb: communication
to invite	verb: communication
to plead	verb: communication
to preach	verb: communication
to read	verb: communication
to request	verb: communication
to say	verb: communication
to scream	verb: communication
to shout	verb: communication
to speak	verb: communication
to suggest	verb: communication
to talk	verb: communication
to teach	verb: communication
to tell	verb: communication
to threaten	verb: communication
to warn	verb: communication
to whine	verb: communication
to whisper	verb: communication
to write	verb: communication
to yell	verb: communication
to build	verb: construction
to construct	verb: construction
to draw	verb: construction
to fix	verb: construction
to make	verb: construction
to paint	verb: construction
to repair	verb: construction
to bump	verb: contact
to crash	verb: contact
to hit	verb: contact
to press	verb: contact
to punch	verb: contact
to slap	verb: contact
to bake	verb: cooking
to boil	verb: cooking
to cook	verb: cooking
to fry	verb: cooking
to grill	verb: cooking
to roast	verb: cooking
to steam	verb: cooking
to bomb	verb: destruction
to break	verb: destruction
to chop	verb: destruction
to destroy	verb: destruction
to kill	verb: destruction
to murder	verb: destruction
to smash	verb: destruction
to stab	verb: destruction
to accept	verb: exchange
to acquire	verb: exchange
to borrow	verb: exchange
to buy	verb: exchange
to donate	verb: exchange
to exchange	verb: exchange
to get	verb: exchange
to give	verb: exchange
to lend	verb: exchange
to loan	verb: exchange
to pay	verb: exchange
to receive	verb: exchange
to sell	verb: exchange
to take	verb: exchange
to trade	verb: exchange
to want	verb: exchange
to burn	verb: light emission
to flame	verb: light-emission
to flicker	verb: light-emission
to glow	verb: light-emission
to shine	verb: light-emission
to sparkle	verb: light-emission
to approach	verb: motion-direction
to arrive	verb: motion-direction
to ascend	verb: motion-direction

to come	verb: motion-direction
to descend	verb: motion-direction
to enter	verb: motion-direction
to escape	verb: motion-direction
to fall	verb: motion-direction
to follow	verb: motion-direction
to go	verb: motion-direction
to lead	verb: motion-direction
to leave	verb: motion-direction
to return	verb: motion-direction
to rise	verb: motion-direction
to sink	verb: motion-direction
to bounce	verb: motion-manner
to chase	verb: motion-manner
to creep	verb: motion-manner
to dive	verb: motion-manner
to drive	verb: motion-manner
to fly	verb: motion-manner
to halt	verb: motion-manner
to hop	verb: motion-manner
to jog	verb: motion-manner
to limp	verb: motion-manner
to march	verb: motion-manner
to pedal	verb: motion-manner
to ride	verb: motion-manner
to run	verb: motion-manner
to skid	verb: motion-manner
to slide	verb: motion-manner
to stagger	verb: motion-manner
to step	verb: motion-manner

to stop	verb: motion-manner
to swerve	verb: motion-manner
to swim	verb: motion-manner
to travel	verb: motion-manner
to wade	verb: motion-manner
to walk	verb: motion-manner
to wander	verb: motion-manner
to chime	verb: noise
to clang	verb: noise
to clatter	verb: noise
to crackle	verb: noise
to rattle	verb: noise
to screech	verb: noise
to sigh	verb: noise
to sing	verb: noise
to snap	verb: noise
to bark	verb: noise-animal
to chirp	verb: noise-animal
to growl	verb: noise-animal
to meow	verb: noise-animal
to oink	verb: noise-animal
to brush	verb: tool-action
to cut	verb: tool-action
to drill	verb: tool-action
to hammer	verb: tool-action
to hoe	verb: tool-action
to pound	verb: tool-action
to rake	verb: tool-action
to saw	verb: tool-action
to shovel	verb: tool-action

Appendix B. Items used in Experiment 1 (semantic priming in lexical decision, objects)

	Primes			
Target	Very close	Close	Medium	Far
apple	peach	lemon	bean	raft
artichoke	carrot	pepper	lime	comb
axe	hammer	spanner	pencil	ceiling
banana	peach	melon	potato	broom
cabbage	onion	pepper	lime	rug
camel	zebra	mouse	swan	sofa
celery	aubergine*	mushroom	watermelon	dustpan
cherry	pear	lemon	bean	spoon
chin	lips	tongue	nose	donkey
coat	suit	shoe	belt	bus
corn	bean	pea	pear	sofa
cucumber	broccoli	pumpkin	strawberry	shield
dagger	sword	razor	hammer	tongue
dog	rabbit	tiger	duck	comb
elbow	wrist	ankle	thumb	tiger
fence	gate	wall	roof	bus
finger	thumb	wrist	knee	couch
hat	scarf	shoe	belt	bomb
hoe	chisel	hatchet	tweezers	tricycle
neck	hair	ear	arm	tail
orange	plum	raisin	pumpkin	dustpan
pig	goat	lion	duck	bomb
pliers	hammer	hatchet	scissors	tricycle
rake	shovel	hatchet	sword	carpet
raspberry	plum	lemon	bean	broom
saw	hammer	drill	pencil	curtain
screwdriver	chisel	hatchet	sword	feather
sheep	goat	zebra	swan	beak
shoulder	arm	leg	thumb	bus
toe	leg	knee	wrist	van
trousers*	shirt	glove	belt	couch
wolf	fox	cow	duck	pen

* Semantic features for this word were obtained for its US English translation equivalent.

Appendix C. Items used in Experiment 2 (semantic priming in lexical decision, actions)

	Primes			
Target	Very close	Close	Medium	Far
ascend	rise	walk	march	write
bake	grill	cook	eat	drop
break	drop	lose	kill	cook
buy	trade	demand	accept	drive
carry	hold	press	stop	look
chat	speak	write	ask	drive
clatter	screech	snore	whine	roast
construct	build	draw	bend	flash
descend	rise	enter	press	speak
dive	swim	wade	boil	knock
drink	swallow	vomit	frown	whine
empty	pour	mix	spray	bark
fill	spray	blend	invite	blink
fix	repair	build	destroy	touch
hiccup	cough	sigh	yell	shave
hop	run	rise	leave	trade
kick	walk	stand	rise	build
lend	trade	demand	accept	touch
lift	hold	press	fall	write
listen	hear	sing	call	hit
oink	chirp	clang	snore	squint
plead	demand	accept	suggest	drive
preach	suggest	argue	smoke	steam
punch	slap	stab	ache	yawn
read	write	speak	ask	eat
scream	yell	chirp	snap	eject
sell	borrow	acquire	invite	sparkle
shout	yell	clang	rattle	itch
sneeze	breathe	smell	vomit	shake
stir	twist	bounce	wander	flame
taste	eat	cook	spit	pull
teach	advise	suggest	request	swallow

Appendix D. Items used in Experiment 3 (picture-word interference, objects)

	Distracters			
Target	Very close	Close	Medium	Far
apple	peach	lemon	bean	raft
axe	hammer	spanner	pencil	ceiling
banana	peach	melon	potato	broom
camel	zebra	mouse	swan	sofa
celery	aubergine*	mushroom	watermelon	dustpan
cherry	pear	lemon	bean	spoon
coat	suit	shoe	belt	bus
corn	bean	pea	pear	sofa
cucumber	broccoli	pumpkin	strawberry	shield
dog	rabbit	tiger	duck	comb
fence	gate	wall	roof	bus
finger	thumb	wrist	knee	couch
hand	arm	leg	thumb	bus
hat	scarf	shoe	belt	bomb
hoe	chisel	hatchet	tweezers	tricycle
lettuce	onion	pepper	lime	rug
orange	plum	raisin	pumpkin	dustpan
pig	goat	lion	duck	bomb
pliers	hammer	hatchet	scissors	tricycle
rake	shovel	hatchet	sword	carpet
saw	hammer	drill	pencil	curtain
screwdriver	chisel	hatchet	sword	feather
sheep	goat	zebra	swan	beak
trousers*	shirt	glove	belt	couch

* Semantic features for this word were obtained for its US English translation equivalent.

Appendix E. Items used in Experiment 4 (picture-word interference, actions)

	Distracters		
Target	Very close	Medium	Far
bleed	ache	stab	sigh
bounce	shake	push	wash
cough	hiccup	snore	squint
dive	swim	pour	glow
drill	build	repair	smile
drink	swallow	grin	warn
eat	taste	vomit	kill
hop	step	rise	buy
kick	run	stand	trade
knock	rattle	growl	shine
pound	slap	bump	lick
press	hold	carry	speak
slide	push	drag	argue
sneeze	breathe	smell	borrow
stop	enter	fall	touch
talk	call	read	follow
throw	hold	leave	call
wade	swim	steam	twist
walk	jog	wander	donate
bleed	ache	stab	sigh
bounce	shake	push	wash
cough	hiccup	snore	squint
dive	swim	pour	glow
drill	build	repair	smile
drink	swallow	grin	warn
eat	taste	vomit	kill
hop	step	rise	buy
kick	run	stand	trade
knock	rattle	growl	shine
pound	slap	bump	lick
press	hold	carry	speak
slide	push	drag	argue
sneeze	breathe	smell	borrow
stop	enter	fall	touch
talk	call	read	follow
throw	hold	leave	call
wade	swim	steam	twist
walk	jog	wander	donate

Appendix F. Items used in Experiment 5 (semantic blocking in picture naming)

Objects:

Body-parts: arm; finger; foot; hand; leg; neck; shoulder; thumb

Clothing: belt; glove; hat; shirt; shoe; sock; trousers*; waistcoat*

Vehicles: aeroplane*; bicycle; bus; car; helicopter; lorry*; motorcycle; train

Actions:

Body actions: hop; kick; march; run; sit; slide; stop; walk

Mouth actions: drink; eat; frown; smile; sneeze; spit; taste; yawn

Tool actions: cut; dig; draw; drill; paint; rake; saw; shovel

* Semantic features for this word were obtained for its US English translation equivalent.

References

Aguirre, E. and Rigau, G. (1996). Word sense disambiguation using conceptual density. In Proceedings of the 16th International Conference on Computational Linguistics, Copenhagen.

Allport, P. (1985). Distributed memory, modular subsystems and dysphasia. In S. K. N. R. Epstein (Ed.) Current perspectives in dysphasia. (pp. 32-60). Edinburgh, Churchill Livingstone.

Anderson, J. (1983). The architecture of cognition. Cambridge, MA: Harvard University Press.

Andrews, M., Vigliocco, G. & Vinson, D. (2005a). Integrating attributional and distributional information in a probabilistic model of meaning representation. In Timo Honkela, Ville Könönen, Matti Pöllä, and Olli Simula, editors, Proceedings of AKRR'05, International and Interdisciplinary Conference on Adaptive Knowledge Representation and Reasoning. Pages 15-25. Espoo, Finland, June 2005.

Andrews, M., Vigliocco, G. & Vinson, D.P. (submitted, revision invited). Integrating attributional and distributional data to learn semantic representations. Psychological Review

Andrews, M., Vigliocco, G., & Vinson, D.P (2005b). The role of attributional and distributional information in representing meaning. Proceedings of the 27th Meeting of the Cognitive Science Society.

Aristotle (350BCE/1941), Categoriae [Categories] (E.M. Edghill, trans). In Richard McKeon (ed.), The Basic Works of Aristotle, pp.7-37. New York: Random House.

Ballard, D.H., Hayhoe, M.M., Pook, P.K. & Rao, R.P.N. (1997). Deictic codes for the embodiment of cognition. Behavioral and Brain Sciences, **20**, 723-767.

Barclay,]., Bransford, J., Franks, J., McCarrell, N. & Nitsch, K. (1974). Comprehension and semantic flexibility. Journal of Verbal Learning and Verbal Behavior, **13**, 471-481.

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

Barsalou, L. W., Kyle Simmons, W., Barbey, A. K., & Wilson, C. D. (2003). Grounding conceptual knowledge in modality-specific systems. Trends in Cognitive Sciences, **7**, 84-91.

Barsalou, L.W., & Wiemer-Hastings, K. (2005). Situating abstract concepts. In D. Pecher and R. Zwaan (Eds.), Grounding cognition: The role of perception and action in memory, language, and thought (pp. 129-163) . New York: Cambridge University Press.

Basso, A., Capitani, E., Laiacona, M. (1988). Progressive language impairment without dementia: a case with isolated category-specific semantic defect. Journal of Neurology, Neurosurgery and Psychiatry, **51**, 1201-1207.

Bates, E.A., & MacWhinney, B. (1982). Functionalist approaches to grammar. In E. Wanner & L.R. Gleitman (Eds.), Language acquisition: The state of the art, (pp.173-218.) Cambridge: Cambridge University Press.

Berlin, B. & Kay, P. (1969). Basic color terms: Their universality and evolution, Berkeley: University of California Press.

Bierwisch, Manfred. (1969). Strukturelle Semantik. Deutsch als Fremdsprache, **6**, 66-71.

Bird, H., Howard, D., & Franklin, S. (2000). Why is a verb like an inanimate object? Grammatical category and semantic category deficits. Brain and Language, **72**, 246–309.

Bird, H., Lambon Ralph, M.A., Patterson, K. & Hodges, J. R. (2000). The rise and fall of frequency and imageability: Noun and verb production in semantic dementia. Brain and Language, **73**, 17-49.

Borgo, F. and Shallice, T. (2001). When living things and other 'sensory-quality' categories behave in the same fashion: a novel category-specific effect. Neurocase, **7**, 201-220.

Borreggine, K. L., & Kaschak, M. P. (2006). The action-sentence compatibility effect: It's all in the timing. Cognitive Science, **30**, 1097-1112.

Bourne, L.E. (1970). Knowing and using concepts. Psychological Review, **77**, 546-556.

Bowerman, M. and Choi, S. (2003). Space under construction: Language specific spatial categorization in first language acquisition. In D. Gentner and S. Goldin-Meadow (Eds) Language in Mind: Advances in the study of Language and Cognition (pp. 387-428). Cambridge: MIT Press.

Breedin, S. D., Saffran, E.M. & Schwartz, M.F. (1998). Semantic factors in verb retrieval: An effect of complexity. Brain and Language, **63**, 1-31.

Breedin, S.D., Saffran, E.M. & Coslett, H.B. (1994). Reversal of the concreteness effect in a patient with semantic dementia. Cognitive Neuropsychology, **11**, 617-660.

Brysbaert, M. (1995). Arabic number reading: On the nature of the numeric scale and the origin of phonological recoding. Journal of Experimental Psychology: General, **124**, 434-452.

Brysbaert, M., Fias, W, & Noel, M.P. (1998). The Whorfian hypothesis and numerical cognition: Is 'twenty-four' processed in the same way as 'four-and-twenty'? Cognition, **66**, 51-77.

Buccino, G., Riggio, L., Melli, G., Binkofski, F., Gallese, V., & Rizzolatti, G. (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.

Budanitsky, A. & Hirst, G. (2001). Semantic distance in Wordnet: An experimental, application-oriented evaluation of five measures. In Proceedings of the North American Chapter of the Association for Computational Linguistics, Pittsburgh.

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

Bushell, C.M., & Martin, A. (1997). Automatic semantic priming of nouns and verbs in patients with Alzheimer disease. Neuropsychologia, **35**, 1059-1067.

Buxbaum, L.J., Veramonti, T. & Schwartz, M.F. (2000). Function and manipulation tool knowledge in apraxia: knowing 'what for' but not 'how'. Neurocase, **6**, 83-97.

Caramazza, A. & Shelton, J. R. (1998). Domain-specific knowledge systems in the brain: The animate-inanimate distinction. Journal of Cognitive Neuroscience, **10**, 1-34.

Casasanto, D. & Lozano, S. (2007). Embodied Language Production: Evidence from gesture, speech disfluency, andmotor action. Embodied Sentence Processing: behavioural, neuropsychological, and computational perspectives. Saarland, Germany.

Cassirer, E. (1953). The Philosophy of Symbolic Forms. Volume 1: Language. (R. Manheim, trans.). New Haven: Yale University Press.

Chiarello, C., Liu, S., Shears, C., & Kacinik, N. (2002). Differential asymmetries for recognizing nouns and verbs: Where are they? Neuropsychology, **16**, 35-48.

Collins, A. C. & Loftus, E.F. (1975). A spreading-activation theory of semantic processing. Psychological Review, **82**, 407-428.

Collins, A. M. & Quillian, M.R. (1969). Retrieval time from semantic memory. Journal of Verbal Learning and Verbal Behavior, **12**, 240-247.

Cree, G. S., McNorgan, C., & McRae, K. (2006). Distinctive features hold a privileged status in the computation of word meaning: Implications for theories of semantic memory. Journal of Experimental Psychology: Learning, Memory, & Cognition, **32**, 643-658.

Cree, G. S., McRae, K. & McNorgan, C. (1999). An attractor model of lexical conceptual processing: Simulating semantic priming. Cognitive Science, **23**, 371-414.

Cree, G.S. & McRae, K. (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.

Crosson, B., Moberg, P.J., Boone, J.R., Gonzalez Rothi, L.J. & Raymer, A. (1997). Category-specific naming deficit for medical terms after dominant thalamic/capsular hemorrhage. Brain and Language, **60**, 407-442.

Damasio, H., Tranel, D., Grabowski, T. J., Adolphs, R., & Damasio, A. R. (2004). Neural systems behind word and concept retrieval. Cognition, **92**, 179-229.

Damian, M. F., Vigliocco, G., & Levelt, W. J. M. (2001). Effects of semantic context in the naming of pictures and words. Cognition, **81**, B77-B86.

Davidoff, J., Davies, I., & Roberson, D. (1999). Colour categories in a stone-age tribe. Nature, **398**, 203-204.

Dell, G.S. (1995). Speaking and misspeaking. In L. Gleitman & M. Liberman (Eds.), Invitation to Cognitive Science, Part I, Language. Cambridge, MA: MIT Press.

Devlin, J. T., Gonnerman, L. M., Andersen, E. S. & Seidenberg, M. S. (1998). Category-specific semantic deficits in focal and widespread brain damage: A computational account. Journal of Cognitive Neuroscience, **10**, 77-94.

Druks, J., & Masterson, J. (2000). An object and action naming battery. London: Psychology Press.

Elman, J. (2003). Development: It's about time. Developmental Science, **6**, 430-433.

Farah, M. J. & McClelland, J.L. (1991). A computational model of semantic memory impairment: Modality specificity and emergent category specificity. Journal of Experimental Psychology: General, **120**, 339-357.

Farah, M.J., Hammond, K.M., Mehta, Z., Ratcliff, G. (1989). Category-specificity and modality-specificity in semantic memory. Neuropsychologia, **7**, 193-200.

Fellbaum, C. (1990). English verbs as a semantic net. International Journal of Lexicography, **3**, 278 - 301.

Fromkin, V. (Ed.), (1973). Speech errors as linguistic evidence. The Hague: Mouton.

Gallese, V., & Lakoff, G. (2005). The brain's concepts: The role of the sensory-motor system in conceptual knowledge. Cognitive Neuropsychology, **22**, 455.

Garnham, A., Shillock, R., Brown, G., Mill, A. and Cutler, A. (1981). Slips of the tongue in the London-Lund corpus of spontaneous conversation. Linguistics. **19**, 805-817.

Garrard, P., Lambon Ralph, M.A., Hodges, J.R. & Patterson, K. (2001). Prototypicality, distinctiveness, and intercorrelation: Analyses of the semantic attributes of living and nonliving concepts. Cognitive Neuropsychology, **18**, 125-174.

Garrett, M. F. (1992). Lexical retrieval processes: Semantic field effects. In: Frames, fields and contrasts: New essays in semantic and lexical organization. E. Kittay & A. Lehrer (Eds.). Hillsdale, NJ, Erlbaum.

Glaser, W. R., & Düngelhoff, F. J. (1984). The time course of picture word interference. Journal of Experimental Psychology: Human Perception and Performance, **10**, 640-654.

Gleitman, L. R. (1990). The structural sources of verb meanings. Language Acquisition, **1**, 3–55.

Gleitman, L., Cassidy, K., Nappa, R., Papafragou, A., & Trueswell, J. (2005). Hard words. Language Learning and Development, **1**, 23-64.

Glenberg, A. M., & Kaschak, M. P. (2002). Grounding language in action. Psychonomic Bulletin and Review, **3**, 558-565.

Glenberg, A. M., & Kaschak, M. P. (2003). The body's contribution to language. In B.H.Ross (Ed.), The Psychology of Learning and Motivation (Vol. 43, pp. 93-126). San Diego, CA: Academic Press.

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**, 379-401.

Gonnerman, L.M., Seidenberg, M.S. & Andersen, E.S. (2007). Graded semantic and phonological similarity effects in priming: Evidence for a distributed connectionist approach to morphology. Journal of Experimental Psychology: General, **136**, 323-345.

Graesser, A. C., Hopkinson, P. L. & Schmid, C. (1987). Differences in interconcept organization between nouns and verbs. Journal of Memory and Language, **26**, 242-253.

Grimshaw, J. (1991). Argument structure. London, Massachusetts Institute of Technology Press.

Gross, D.; Fischer, U. & Miller, G. A. (1989). Antonymy and the representation of adjectival meanings. Journal of Memory and Language, **28**, 92-106.

Hampton, J.A. (1979). Polymorphous concepts in semantic memory. Journal of Verbal Learning and Verbal Behavior, **18**, 441-461.

Hampton, J.A. (1981). An investigation of the nature of abstract concepts. Memory & Cognition, **9**, 149-156.

Hampton, J.A., & Gardiner, M.M. (1983). Measures of internal category structure: A correlational analysis of normative data. British Journal of Psychology, **74**, 491-516.

Hart, J., Berndt, R.S., & Caramazza, A. (1985). Category-specific naming deficit following cerebral infarction. Nature, **316**, 439-440.

Hauk, O., Johnsrude, I., & Pulvermüller, F. (2004). Somatotopic representation of action words in human motor and premotor cortex. Neuron, **41**, 301-307.

Hillis, A. E. & Caramazza, A. (1991). Category-specific naming and comprehension impairment - a double dissociation. Brain, **114**, 2081-2094.

Hillis, A. E. & Caramazza, A. (1995). Representation of grammatical categories of words in the brain. Journal of Cognitive Neuroscience, **7**, 396-407.

Hinton, G.E. & Shallice, T. (1991). Lesioning an attractor network: Investigations of acquired dyslexia. Psychological Review, **98**, 74-95.

Humphreys, G.W., Price, C.J., & Riddoch, M.J. (1999). From objects to names: A cognitive neuroscience approach. Psychological Research, **62**, 118-130.

Hutchison, K.A. (2003). Is semantic priming due to association strength or featural overlap? A micro-analytic review. Psychonomic Bulletin & Review. 10, 785–813.

Huttenlocher, J. & Lui, F. (1979). The semantic organization of some simple nouns and verbs. Journal of Verbal Learning and Verbal Behavior, **18**, 141-179.

Jackendoff, R. (1990). Semantic structures. Cambridge, MA, MIT Press.

Jackendoff, R. (1992). Languages of the mind. Cambridge, MA: MIT Press

Jackendoff, R. (2002). Foundations of language. Oxford: Oxford University Press.

Jesperson, O. (1965). A modern English grammar based on historical principles. London, Allen & Unwin.

Johnson-Laird, P.N., Herrmann, D.J., and Chaffin, R. (1984) Only connections: a critique of semantic networks. Psychological Bulletin, **96**, 292-315.

Katz, J. J. & Fodor, J. A. (1963) The structure of a semantic theory, Language, **39**, 170-210.

Keil, F. C. (1989). Concepts, kinds, and cognitive development. Cambridge, MA, MIT Press.

Kittay, E. F. (1987). Metaphor: Its cognitive force and linguistic structure. Oxford, Clevender Press.

Klopfer, D. S. (1996). Stroop interference and color-word similarity. Psychological Science, **7**, 150-157.

Kohonen, T. (1997). Self-organizing maps, New York: Springer.

Kousta, S-T., Vinson, D.P. & Vigliocco, G. (in press). Investigating linguistic relativity through bilingualism: The case of grammatical gender. Journal of Experimental Psychology: Learning, Memory and Cognition.

Kousta, S-T., Vinson, D.P., Andrews, M. & Vigliocco, G. (submitted). The representation of abstract word meanings: Why emotion matters.

Kroll, J. F., & Stewart, E. (1994). Category interference in translation and picture naming: evidence for asymmetric connections between bilingual memory representations. Journal of Memory and Language, **33**, 149-174.

Kucera and Francis, W.N. (1967). Computational analysis of present-day American English. Providence: Brown University Press.

Lakoff, G. (1987). Women, fire, and dangerous things. What categories reveal about the mind. Chicago/London, The University of Chicago Press.

Landau, B. & Gleitman, L. (1985). Language and experience: Evidence from the blind child. Cambridge, MA: Harvard University Press

Landauer, T.K. & Dumais, S.T. (1997). A solution to Plato's problem: The Latent Semantic Analysis theory of acquisition, induction and representation of knowledge. Psychological Review, **104**, 211-240.

Lehrer, A. (1974). Semantic fields and lexical structure. Amsterdam: North-Holland Publishing.

Levelt, W. J. L. (1989). Speaking: From intention to articulation. Cambridge, MA, MIT Press.

Levelt, W. J. M., Roelofs, A., & Meyer, A. S. (1999). A theory of lexical access in speech production. Behavioral and Brain Sciences, **22**,1-38.

Levin, B. (1993). English verb classes and alternations. A preliminary investigation. Chicago, University of Chicago Press.

Levinson, S.C. (1996). Frames of reference and Molyneux's question: Crosslinguistic evidence. In P. Bloom, M. Peterson, L. Nadel, & M. Garrett (Eds.), Language and space, pp. 109-169. Cambridge, MA: MIT Press.

Lewis, W. (2002). Measuring conceptual distance using WordNet: The design of a metric for measuring semantic similarity. In R. Hayes, W. Lewis, E. Obryan, and T. Zamuner (Eds.), The University of Arizona working papers in linguistics. Tucson: University of Arizona.

Lucas, M. (2000). Semantic priming without association: A meta-analytic review. Psychonomic Bulletin & Review, **7**, 618-630.

Lucy, J.A. (1992). Grammatical categories and cognition: A case study of the Linguistic Relativity Hypothesis. Cambridge, UK: Cambridge University Press.

Lupker, S. J. (1979). The semantic nature of response competition in the picture–word interference task. Memory & Cognition, **7**, 485-495.

Maki, W.S., Krimsky, M. & Muñoz, S. (2006). An efficient method for estimating semantic similarity based on feature overlap: Reliability and validity of semantic feature ratings. Behavior Research Methods, **38**, 153-157.

Malt, B. C. Sloman, S.A., Gennari, S., Shi, M. & Wang, Y. (1999). Knowing versus naming: Similarity and the linguistic categorization of artifacts. Journal of Memory and Language, **40**, 230-262.

Martin, A., & Chao, L. L. (2001). Semantic memory and the brain: structure and processes. Current Opinion in Neurobiology, **11**, 194-201.

McKenna, P. and Warrington, E. K. (1978). Category-specific naming preservation: A single case study. Journal of Neurology, Neurosurgery, and Psychiatry, **41**, 571-574.

McRae, K., & Boisvert, S. (1998). Automatic semantic similarity priming. Journal of Experimental Psychology: Learning, Memory and Cognition, **24**, 558-572.

McRae, K., & Cree (2002) Factors underlying category-specific semantic deficits. In E. Forde & G. Humphreys (Eds.) Category specificity in brain and mind, Psychology Press.

McRae, K., Cree, G. S., Seidenberg, M. S., & McNorgan, C. (2005). Semantic feature production norms for a large set of living and nonliving things. Behavioral Research Methods, Instruments, and Computers, **37**, 547-559

McRae, K., Cree, G. S., Westmacott, R., & de Sa, V. R. (1999). Further evidence for feature correlations in semantic memory. Canadian Journal of Experimental Psychology: Special Issue on Models of Word Recognition, **53**, 360-373.

McRae, K., de Sa, V. & Seidenberg, M. (1997). On the nature and scope of featural representations of word meaning. Journal of Experimental Psychology: General, **126**, 99-130.

Meng, X.L, Rosenthal, R. & Rubin, D.B. (1992). Comparing correlated correlation coefficients. Psychological Bulletin, **111**, 172-175.

Mervis, C.B. & Rosch, E. (1981). Categorization of natural objects. Annual Review of Psychology, **32**, 89-115.

Meteyard, L. & Vigliocco, G. (in press) The role of Sensory and Motor Information in Semantic Representation: a review. Elsevier Handbook of Embodied Cognitive Science.

Meteyard, L. (2008) Motion seen and understood: Interactions between language comprehension and visual perception. Unpublished PhD thesis, University College London.

Meteyard, L., Bahrami, B. & Vigliocco, G. (2007) Motion detection and motion words: language affects low level visual perception. Psychological Science, **18**, 1007-1013.

Meteyard, L., Zokaei, N., Bahrami, B. & Vigliocco, G. (in press) Now you see it: visual motion interferes with lexical decision on motion words. Current Biology.

Meyer, D.E. & Schvaneveldt, R.W. (1971). Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. Journal of Experimental Psychology, **90**, 227-234.

Miller, G. A. & Fellbaum, C. (1991). Semantic networks of English. Cognition, **41**, 197-229.

Miller, G. A. (1995). Wordnet: A lexical database for English. Communications of the Association for Computing Machinery, **38**, 39-41.

Miller, G. A., & Johnson-Laird, P. N. (1976). Language and perception. Cambridge, MA: Belknap Press.

Minsky, M. (1975). A framework for representing knowledge. In: P. H. Winston (Ed.), The psychology of computer vision. (pp. 211-277). New York, McGraw-Hill.

Morris, M. W. & Murphy, G.L. (1990). Converging operations on a basic level in event taxonomies. Memory and Cognition, **18**, 407-418.

Moss, H. E., & Tyler, L. K. (2000) A progressive category-specific deficit for non-living things. Neuropsycholgia, **38**, 60-82.

Moyer, R. S., & Landauer, T. K. (1967). Time required for judgements of numerical inequality. Nature, **215**, 1519-1520.

Murphy, G.L. (2002). The big book of concepts. Cambridge, MA: MIT Press.

Najmi, S. & Wegner, D.M. (2008). The gravity of unwanted thoughts: Asymmetric priming effects in thought suppression. Consciousness and Cognition, **17**, 114-124.

Neely, J. H. (1991). Semantic priming effects in visual word recognition: A selective review of current findings and theories. In D. Laberge & S. J. Samuels (Eds.), Basic processes in reading: Perception and comprehension. Hillsdale, NJ, Erlbaum.

Norman, D. A. & Rumelhart, D. E. (1975). Explorations in cognition. San Francisco, Freeman.

Osgood, C. E., May, W. H., and Miron, M. S. (1975) Cross-cultural universals of affective meaning. Urbana, IL: University of Illinois Press.

Osgood, C. E., Suci, G. J., & Tannenbaum, P. H. (1957). The measurement of meaning. Urbana: University of Illinois Press.

Osgood, C.H. (1962). Studies of the generality of affective meaning systems. American Psychologist **17**, 10-28.

Paivio, A. (1971). Imagery and verbal processes. New York: Holt, Rinehart, and Winston.

Paivio, A. (1986). Mental representations: a dual coding approach. Oxford: Oxford University Press.

Paivio, A. (1991). Images in mind: The evolution of a theory. New York: Harvester Wheatsheaf.

Paivio, A. (2007). Mind and its evolution: A dual coding theoretical approach. New York: Lawrence Erlbaum.

Pavesi, A., & Umilta`, C. (1998). Symbolic distance between numerosity and identity modulates Stroop interference. Journal of Experimental Psychology: Human Perception and Performance, **24**, 1535-1545.

Pinker, S. (1989). Learnability and cognition: The acquisition of argument structure. Cambridge, MA, MIT Press.

Plaut, D. C. & Shallice, T. (1993). Deep dyslexia: A case study of connectionist neuropsychology. Cognitive Neuropsychology, **10**, 377-500.

Plaut, D. C. (1995). Double dissociation without modularity: evidence from connectionist neuropsychology. Journal of Clinical and Experimental Neuropsychology, **17**, 291-321.

Pulvermüller, F. (2001). Brain reflections of words and their meaning. Trends in Cognitive Sciences, **5**, 517-524.

Randall, B., Moss, H. E., Rodd, J. M., Greer, M., & Tyler, L. K. (2004). Distinctiveness and correlation in conceptual structure: Behavioral and computational studies. Journal of Experimental Psychology: Learning, Memory, and Cognition, **30**, 393-406.

Richardson, R., Smeaton, A.F., & Murphy, J. (1994). Using WordNet as a knowledge base for measuring semantic similarity between words. Technical Report, Working paper CA-1294, School of Computer Applications, Dublin City University: Dublin, Ireland.

Roberson, D., Davies I. & Davidoff, J. (2000) Color categories are not universal: Replications and new evidence from a Stone-age culture. Journal of Experimental Psychology: General, **129**, 369-398.

Roelofs, A. (1993). Testing a non-decompositional theory of word retrieval in speaking: retrieval of verbs. Cognition, **47**, 59-87.

Rogers, T. T. and Plaut, D. C. (2002) Connectionist perspectives on category-specific deficits. In E. Forde and G. Humphreys (eds.), Category specificity in brain and mind. Brighton: Psychology Press, 251-289.

Rogers, T. T., and J. L. McClelland (2004). Semantic Cognition: A Parallel Distributed Processing Approach. Cambridge, MA: MIT Press.

Rogers, T. T., Lambon Ralph, M. A, Garrard, P., Bozeat, S., McClelland, J. L., Hodges, J. R., and Patterson, K. (2004). The structure and deterioration of semantic memory: A neuropsychological and computational investigation. Psychological Review, **111**, 205-235.

Rosch, E. & Mervis, C.B. (1975). Family resemblance: Studies in the internal structure of categories. Cognitive Psychology, **7**, 573-605.

Rosch, E.H. (1973) Natural categories, Cognitive Psychology, **4**, 328-50.

Rosch, E.H., Mervis, C.B., Gray, W.D., Johnson, D.M. and Boyes-Braem, P. (1976) Basic objects in natural categories, Cognitive Psychology, **8**, 382-439.

Rösler, F., Streb, J.,& Haan, H. (2001). Event-related brain potentials evoked by verbs and nouns in a primed lexical decision task. Psychophysiology, **38**, 694-703.

Sacchett, C. & Humphreys, G. W. (1992). Calling a squirrel a squirrel and a canoe a wigwam: A category-specific deficit for artifactuals and body parts. Cognitive Neuropsychology, **9**, 73-86.

Sapir, E. (1921). Language, New York, NY: Harcourt, Brace, and World.

Sartori, G. & Job, R. (1988). The oyster with 4 legs: A neuropsychological study of the interaction of visual and semantic information. Cognitive Neuropsychology, **5**, 105-132.

Schneider, W., Eschman, A., & Zuccolotto, A. (2002). E-Prime reference guide. Pittsburgh, PA: Psychology Software Tools.

Schriefers, H., Meyer, A. & Levelt, W.J.M. (1990). Exploring the time course of lexical access in language production: picture-word interference studies, Journal of Memory and Language, **29**, 86-102.

Sera, M., Elieff, C., Forbes, J., Burch, M.C., Rodriguez, W., Dubois, D.P. (2002). When language affects cognition and when it does not: An analysis of grammatical gender and classification. Journal of Experimental Psychology: General, **131**, 377-397.

Shallice, T. (1993). Multiple semantics: Whose confusion? Cognitive Neuropsychology, **10**, 251-261.

Shelton, J. R., Fouch, E., & Caramazza, A. (1998). The selective sparing of body-part knowledge: a case study. Neurocase, **4**, 339-351.

Sheridan, J. & Humphreys, G.W. (1993). A verbal-semantic category-specific recognition impairment. Cognitive Neuropsychology, **10**, 143-184

Siri S., Kensinger E.A., Cappa S.F., Hood K.L., & Corkin S. (2003). Questioning the living- nonliving dichotomy: Evidence from a patient with an unusual semantic dissociation. Neuropsychology, **17**, 630-645.

Slobin, D.I. (1996b). Two ways to travel: Verbs of motion in English and Spanish. In M. S. Shibatani & S. A. Thompson (Eds.), Grammatical constructions: Their form and meaning (pp. 195-220). Oxford: Clarendon Press.

Slobin, D.I. (1996a). From "thought and language" to "thinking for speaking". In J. Gumperz & S. Levinson (Eds.), Rethinking Linguistic Relativity. Cambridge, MA: Cambridge University Press, 70-96.

Smith, E. E. & Medin, D. L. (1981). Categories and concepts. Cambridge, MA, Harvard University Press.

Smith, E. E., Shoben, E.J. & Rips, L.J. (1974). Structure and process in semantic memory: Featural model for semantic decisions. Psychological Review, **81**, 214-241.

Snider, J. G., and Osgood, C. E. (1969), Semantic Differential technique: A sourcebook. Chicago: Aldine

Snodgrass, J. G., & Vanderwart, M. (1980). A standardized set of 260 pictures: norms for name agreement, image agreement, familiarity, and visual complexity. Journal of Experimental Psychology: Learning, Memory and Cognition, **6**, 174-215.

Stroop, J.R. (1935). Studies of interference in serial verbal reactions. Journal of Experimental Psychology, **18**, 643-662.

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**, 273-281.

Trier, J. (1931) Der Deutsche Wortschatz im Sinnbezirk des Verstandes. Heidelberg: Winter.

Trueswell, J., & Gleitman, L. (2004). Children's eye movements during listening: Developmental evidence for a constraint-based theory of sentence processing. In J. M. Henderson & F. Ferreira (Eds.), The interface of language, vision, and action: Eye movements and the visual world (pp. 319–346). New York: Taylor & Francis.

Tversky, B. and Hemenway, K. (1984). Objects, parts, and categories. Journal of Experimental Psychology: General, **113**, 169-193.

Tyler, L. K., Moss, H.E., Durrant-Peatfield, M.R. & Levy, J.P. (2000). Conceptual structure and the structure of concepts: A distributed account of category-specific deficits. Brain and Language, **75**, 195-231.

Vigliocco, G., Barber, H., Vinson, D.P., Druks, J. & Cappa, S.F. (in prep). Nouns and verbs in the brain: A Review of behavioural, electrophysiological, neuropsychological and imaging studies.

Vigliocco, G. & Filipovic, L. (2004). From mind in the mouth to language in the mind. Trends in Cognitive Science, **8**, 5-7.

Vigliocco, G. & Vinson, D.P. (2007). Semantics. In G. Gaskell (Ed.) Oxford Handbook of Psycholinguistics. Oxford University Press.

Vigliocco, G., Vinson, D.P, Lewis, W. & Garrett, M.F. (2004). The meanings of object and action words. Cognitive Psychology **48**, 422-488.

Vigliocco, G., Vinson, D.P., Arciuli, J. & Barber, H. (2008). The role of grammatical class on word recognition. Brain and Language, **105**, 175-184.

Vigliocco, G., Vinson, D.P., Damian, M.F. & Levelt, W. (2002). Semantic distance effects on object and action naming. Cognition, **85**, B61-B69.

Vigliocco, G., Vinson, D.P., Paganelli F. & Dworzynski, K. (2005). Grammatical gender effects on cognition: Implications for language learning and language use. Journal of Experimental Psychology: General, **134**, 501-520.

Vigliocco, G., Vinson, D.P., Siri, S. (2005). Semantic and grammatical class effects in naming actions. Cognition, **94**, B91-B100.

Vigliocco, G., Warren, J., Siri, S., Arcuili, J., Scott, S.K., Wise, R. (2006). The role of semantics and grammatical class in the neural representation of words. Cerebral Cortex, **16**, 1790-1796.

Vigliocco,G., Vinson,D.P., Woolfe,T., Dye,M., Woll,B. (2005). Language and Imagery: effects of language modality. Proceedings of the Royal Society B, **272**, 1859-1863.

Vinson, D. P., & Vigliocco, G. (2002). A semantic analysis of grammatical class impairments: semantic representations of object nouns, action nouns, and action verbs. Journal of Neurolinguistics, **15**, 317-351.

Vinson, D.P., & Vigliocco, G. (2008). Semantic feature production norms for a large set of objects and events. Behavior Research Methods, **40**, 183-190.

Vinson, D.P., Vigliocco, G., Cappa, S.F. & Siri, S. (2003). The breakdown of semantic knowledge: Insights from a statistical model of meaning representation. Brain and Language, **86**, 347-442.

Vitkovitch, M., Humphreys, G. W., & Lloyd-Jones, T. (1993). On naming a giraffe a zebra: picture naming errors across different categories. Journal of Experimental Psychology: Learning, Memory and Cognition, **19**, 243-259.

Warrington, E. K. & McCarthy, R.A. (1987). Categories of knowledge: Further fractionations and an attempted integration. Brain, **110**, 1273-1296.

Warrington, E. K. & Shallice, T. (1984). Category specific semantic impairments. Brain, **107**, 829-854.

Whorf, B. (1956). Language, thought, and reality: selected writing of Benjamin Lee Whorf, ed. J.B. Carroll. Cambridge, MA: MIT Press.

Wickens, D. D. (1970). Encoding categories of words: an empirical approach to meaning. Psychological Review, **77**, 1-15.

Wickens, D. D., Born, D. G., & Allen, C. K. (1963). Proactive inhibition and item similarity in short-term memory. Journal of Verbal Learning and Verbal Behavior, **2**, 440-445.

Wittgenstein, L. (1953/2001). Philosophical Investigations. Blackwell Publishing.

Zwaan, R., & Taylor, L. J. (2006). Seeing, acting, understanding: Motor resonance in language comprehension. Journal of Experimental Psychology: General, **135**, 1-11.