Conceptual Structure:

Helen E Moss, Lorraine K Tyler & Kirsten I Taylor

Chapter for Oxford Handbook of Psycholinguistics

Address for Correspondence
Lorraine K Tyler
Centre for Speech and Language
Department of Experimental Psychology
University of Cambridge
Downing Street
Cambridge CB2 3EB

lktyler@csl.psychol.cam.ac.uk
Hem10@cam.ac.uk

Running Head: Conceptual Structure

1. **INTRODUCTION**

Concepts lie at the heart of our mental life, supporting cognitive functions from language comprehension and production to reasoning, remembering and recognising objects. Therefore, the study of the representation and processing of conceptual knowledge has been a central activity across many disciplines, traditionally in the realm of philosophy, and more recently psycholinguistics, neuropsychology and neuroscience. Given the context of the present Handbook, we will focus here on the role of concepts – and conceptual structure – in the comprehension and production of language, although we will consider evidence from beyond the traditional boundaries of psycholinguistics, in particular from cognitive neuropsychology. Brain-damaged patients with deficits in one category or domain of knowledge have provided valuable insights into the nature of conceptual representations, and represent an alternative study population with which to test the claims of psycholinguistically motivated theories of conceptual knowledge. A detailed consideration of relevant neuro-imaging research is beyond the scope of the current chapter, but can be found in Taylor *et al.* (in press). We believe that converging evidence from different approaches and methodologies provides the greatest potential for progressing our understanding of this central but challenging subject.

1.1. CONCEPTS AND MEANINGS

Conceptual representations are essential to express and understand information about the world, so must capture a rich variety of knowledge about objects, abstract ideas, mental states, actions and the relations among all of these. A distinction has frequently been made in the psycholinguistic literature between a much smaller set of definition-like semantic information stored as part of the mental lexicon, and richer "world knowledge" or "encyclopedic" information (e.g. Miller, 1978). In this chapter we do not make such a distinction, assuming rather the position taken by Jackendoff (1983, 1989) in his Conceptual Semantic approach, in which there is no clear lexical/encyclopedic boundary, with both kinds of information being "cut from the same cloth" (Jackendoff, 1983). According to Jackendoff, semantic structures (i.e. the semantic content in the mental lexicon) are simply a subset of conceptual structures which can be verbally expressed. Critically, conceptual representations form an interface between lexical information and other domains such as sensory and motor systems. This function introduces several constraints on the nature of these representations; irrespective of their content, they must be in a format that is readily accessible by both a range of linguistic and non-linguistic modalities of input and output, they must permit processing rapid enough to support on-line production and comprehension of meaningful speech at a rate of several words per second and to support rapid and appropriate motor responses to meaningful sensory stimuli in the environment, and they must enable flexibility of meaning across different contexts of use (see Moss and Gaskell, 1999). For a more detailed discussion and an alternative approach to the relationship between concepts and word meanings, see Vigliocco & Vinson (this volume).

1.2. COMPONENTIALITY

In this chapter we explore the representation of concepts within a theoretical framework that stresses the importance of the internal structure of those representations. We will discuss several specific models within this framework, but present the Conceptual Structure Account (Durrant-Peatfield *et al.*, 1997, Moss *et al.*, 1997a, Moss *et al.*, 2002, Tyler and Moss, 2001, Tyler *et al.*, 2000a, Moss and Tyler, 2000, Taylor *et al.*, in press) in greater detail as one example of this class of accounts. While competing perspectives on conceptual representation can be found in the literature, an extensive discussion of these is not within the purview of this chapter, (see e.g. Barsalou *et al*, 1999; Murphy & Medin, 1985; Vigliocco & Vinson, this volume; Glenberg, this volume).

The notion that concepts *have* an internal structure rests on the critical assumption that conceptual representations are componential in nature; that is, that they are made up of smaller elements of meaning, variously referred to as properties, features or attributes (or in the connectionist literature, as microfeatures). Although agreement on this point has by no means been unanimous, with several theorists arguing that meanings are unanalysable wholes (e.g. Fodor and Fodor, 1980, see also de Almeida, 1999), componentiality is now quite widely assumed in the psycholinguistic literature. Thus, many current models of concept representation (e.g., classical, prototype and exemplar theories) share the assumption that concepts are made up of attributes, and that categorisation is possible to the extent that instances of a concept can be grouped together according to the similarity or overlap of these attributes. Although these models differ with respect to the nature of the attributes they consider and the similarity computations they hypothesise, with classical models considering only necessary and sufficient features while prototype and exemplar models additionally account for characteristic, non-defining features, all share the assumptions of componentiality and similarity (see Smith and Medin, 1981, Komatsu, 1992 for reviews and Viglicco and Vinson, this volume for further discussion of this issue).

The componential approach has gained popularity in recent years, in part due to the rise of parallel distributed processing models of cognitive functions including those of the conceptual system (e.g. Durrant-Peatfield *et al.*, 1997, Hinton *et al.*, 1986, Plaut and Shallice, 1993, Rogers and Plaut, 2002, Rogers *et al.*, 2004, Tyler *et al.*, 2000a). Such models instantiate conceptual knowledge in neural networks in which simple processing nodes correspond to components of meaning, and where individual concepts are captured as patterns of activation over large sets of these microfeatures. Although many of the models were developed to show a "proof of concept" and were thus based on small, arbitrarily chosen feature sets, more recent models have grounded their semantic representations in more realistic, empirically derived datasets (e.g. Greer *et al.*, 2001, McRae *et al.*, 1997, Devlin *et al.*, 1998, see below for further details). These distributed models on the whole appear to account for psycholinguistic and neuropsychological phenomena as well as or better than earlier localist models, in which concepts are represented as a single node, with semantic content captured in activation links among related nodes (e.g. Collins and Loftus, 1975). For example, distributed models can easily perform generalisations to novel stimuli and pattern completions (Hinton *et al.*, 1986), and show a gradual degradation of performance when damaged, reminiscent of that shown following brain damage in humans (e.g. Hinton and Shallice, 1991, Plaut and Shallice, 1993, Devlin *et al.*, 1998, Tyler *et al.*, 2000a, Rogers *et al.*, 2004). Although localist semantic network models have been extremely influential as a framework for interpreting psycholinguistic phenomena (most notably semantic priming, e.g. McNamara, 1992), priming effects can also be readily accommodated within the distributed framework (e.g. Masson, 1995, McRae *et al.*, 1997, Tyler *et al.*, 2002, Cree *et al.*, 1999, Vigliocco *et al.*, 2004, Randall *et al.*, 2004). Moreover, the finer-grained effects of conceptual structure can more readily be accommodated within the distributed framework, including the flexibility of meaning over different contexts (e.g. Tabossi, 1988, Moss and Marslen-Wilson, 1993, Kawamoto, 1993), the effects of the number of semantic features associated with a concept (e.g. Plaut and Shallice, 1993, Pexman *et al.*, 2003, Tyler *et al.*, 2000a) and the varying time courses of activation of different features of a word's meaning (McRae *et al.*, 1997, Randall *et al.*, 2004, Moss *et al.*, 1997a). In summary, a distributed componential model seems to be a well-founded and plausible framework for understanding and studying conceptual representations and their internal structures.

THE CONCEPTUAL STRUCTURE FRAMEWORK

We turn now to the recent literature on conceptual structure, which builds on the basic assumptions of the componential account, aiming to specify the internal structures of various classes of concepts and to determine the processing consequences of these proposed representational characteristics. This is a cyclical process. Given that we cannot observe the content of mental representations directly, we draw inferences about their likely nature on the basis of empirical observations of conceptual processing over a range of tasks and stimuli, from both healthy and impaired language users. These inferences can then be used to construct a hypothetical account of the structure of the underlying concepts, on the basis of which further predictions about the potential processing consequences can be generated – often with the help of a computational model – which in turn can be tested in behavioural studies. Recently it has become possible to test hypotheses concerning the potential neural bases of conceptual representations with neuro-imaging techniques, although the relationship between activations in different areas of the brain and the content and structure of individual concepts remains a controversial one (Caramazza and Mahon, 2003, Moss and Tyler, 2003, Tyler *et al*., 2003, Tyler and Moss, 2001).

2.1. Structure and Features

At the heart of the conceptual structure approach is the claim that a given concept can be defined in terms of the features that make up its meaning, and that the quantity and quality of these features, as well as featural interrelationships, – its internal structure – determines how a concept is activated during normal language comprehension and production, as well as the way that it is affected by damage to the system. Within the general term "structure" we will discuss those variables that have been shown to have the most prominent effects: the total number of features a concept has, whether those features are shared with many other concepts or are highly distinctive, the relationships between those features, most critically the degree of correlation among them, and the kind of information specified by the features, both in terms of their derivation (for example whether they are based on sensory information or are functional in nature) and their salience (whether they are critical or peripheral to the concept). While this list of factors is no doubt incomplete, it captures the main areas of investigation pursued in our own work in this area and that of other researchers with similar interests.

2.2. Structure and Domain

Most models of this type also propose that the internal structure of conceptual representations varies systematically across different categories or domains of knowledge, for example between living and non-living things, abstract and concrete words, or verbs and nouns, a proposal that has been of particular relevance in the neuropsychological literature, where there has been lively debate as to the bases of category-specific semantic deficits. Patients with this type of disorder show dissociations in their degree of impairment for different categories/domains; for example, living and non-living things (Bunn *et al*., 1998, Moss *et al*., 1998, de Renzi and Lucchelli, 1994, Hart and Gordon, 1992, Hillis and Caramazza, 1991, Warrington and Shallice, 1984, see Forde and Humphreys, 1999 for a review), with verbs and nouns (Caramazza and Hillis, 1991, Silveri and Di Betta, 1997, Zingeser and Berndt, 1990, Laiacona and Caramazza, 2004) or concrete and abstract words

(Franklin, 1989, Warrington and Shallice, 1984, Breedin *et al*., 1994).[1] Other patients have demonstrated deficits for specific categories within domains, for example for animals (e.g., EW; Caramazza and Shelton, 1998) or fruit and vegetables (e.g., FAV; Crutch & Warrington, 2003). Some researchers have suggested that these deficits are evidence for the existence of modular neural subsystems for concepts in different domains which can be independently impaired as a result of brain damage (e.g. Caramazza and Shelton, 1998, Caramazza and Mahon, 2003). However, the conceptual structure framework suggests an alternative approach in which category specific semantic deficits do not necessarily result from damage to modular neural sub-systems (for example, selective damage to the animal sub-system). Rather patterns of preserved and impaired conceptual knowledge may emerge within a unitary distributed system – one without explicit boundaries into separate categories or domains. This is possible because the effects of neural damage will vary as a function of the internal structure of concepts, which can be shown to vary systematically across different categories, as will be described in further detail below (Durrant-Peatfield *et al*., 1997, Moss *et al*., 1998, Moss *et al*., 2002, Tyler *et al*., 2000a, Tyler and Moss, 2001, Gonnerman *et al*., 1997, Devlin *et al*., 1998, McRae *et al*., 1997, McRae and Cree, 2002, Garrard *et al*., 2001, Caramazza *et al*., 1990, Cree and McRae, 2003).

2.3. Models in the Conceptual Structure Framework

Several models could be described as belonging within the general framework of the conceptual structure approach, with the earliest forerunner perhaps being the OUCH (organised unitary content hypothesis) model proposed by Caramazza *et al* (1990) as an account of category-specific deficits in neuropsychological patients as an alternative to modular accounts. According to the OUCH model, concepts are represented as patterns of activation over "microfeatures within a distributed semantic system". Damage to this kind of unitary, distributed system can potentially affect one category of concepts more than another because similar concepts are represented close together in semantic space – they have overlapping patterns of activation (see Caramazza *et al*., 1990, Dixon *et al*., 1997, Forde *et al*., 1997, and Humphreys *et al*., 1988 for related similarity-based models).

2.3.1. The Conceptual Structure Account

Our own model, which has acquired the name Conceptual Structure Account (CSA), has been developed over the last decade (Durrant-Peatfield *et al*., 1997, Moss *et al*., 1997a, Moss *et al*., 2002, Tyler and Moss, 2001, Tyler *et al*., 2000a, Moss and Tyler, 2000, Taylor *et al*., in press), building on empirical results from our own studies as well as the findings of other research teams who have proposed similar accounts of conceptual structure, albeit with some critical differences (McRae *et al*., 1997, Gonnerman *et al*., 1997, Devlin *et al*., 1998, Cree and McRae, 2003, Vinson *et al*., 2003). Like Caramazza *et al*'s OUCH model, the CSA was initially proposed to try to account for the nature of category-specific semantic deficits, largely in response to the limitations of modular accounts which hypothesised separate neural subsystems for either conceptual domains or types of featural knowledge. Modular accounts are those which propose that there are topographically distinct regions of the brain underpinning representations of specific domains of knowledge (e.g. animals or plants as in the Caramazza & Shelton (1998) account mentioned above) or specific types of conceptual

[1] There is considerable variation in the literature as to the use of the terms category and domain. Here we use "domain" to refer to a high level grouping of concepts into for example living and non-living things or abstract and concrete words, and reserve the term "category" to denote specific superordinate classes within those domains (e.g. animals, vehicles, motion verbs). However, we acknowledge that these are working definitions, and that there are other ways in which the conceptual space may potentially be subdivided.

feature, such as sensory vs motor features (e.g. Warrington and McCarthy, 1983, Warrington and McCarthy, 1987, Warrington and Shallice, 1984, Borgo and Shallice, 1994, Farah *et al.*, 1989, Farah and McClelland, 1991, Saffran, 2000, and see section 3.1 for further discussion of the representation of different feature types). We present the CSA here as an illustrative example of one model within the theoretical framework of the conceptual structure approach.

On the CSA, not only do similar concepts tend to activate overlapping sets of semantic features, but there are also systematic differences in the *internal structure* of concepts in different categories and domains. Like some other accounts, we claim that *correlation* – the degree to which properties co-occur - is a key relation among semantic properties (see also Garrard *et al.*, 2001, Vinson *et al.*, 2003). In a distributed connectionist system, correlated properties like *has eyes* and *can see* support each other with mutual activation and so are more resilient to damage than are those that are more weakly correlated. Hence, different patterns of correlation among properties within a concept will lead to different patterns of loss and preservation of information, given the same degree of overall damage.

In this framework category-specific deficits arise because the structure of concepts in the living and non-living domains differs in systematic ways. Importantly, the most common type of category specific deficit is an impairment for living things relative to non-living things. Based on well-supported claims in the psychological literature, we proposed that living things (and most typically animals) have many properties and many of these are shared among all members of a category (e.g. all mammals *breathe, move, have eyes, can see, have live young, eat* and so on). Moreover, these shared properties co-occur frequently and so are strongly correlated (Keil, 1986, 1989). Living things also have distinctive properties that are informative in distinguishing one category member from another (e.g. *having stripes vs. having spots*), although these tend to be weakly correlated with other properties and so are vulnerable to damage. Artifacts have fewer properties in total, and they tend to be relatively more distinctive, with a smaller pool of information shared across all members of a category. Moreover, unlike living things, the distinctive properties of artifacts tend to be more highly correlated. The result of these structural differences is that brain damage will tend to adversely affect the distinctive properties of living things (because they are weakly correlated) to a greater extent than those of artifacts (because they are more strongly correlated), so making living things concepts difficult to differentiate, leading to the typical pattern of semantic deficit for living things. On the CSA account, the variables of distinctiveness and intercorrelation also interact with feature *type*, allowing more complex differences across categories to emerge. This is based on the widely held premise that concepts (at least those referring to concrete entities) contain a mix of different types of features, including both perceptually based (e.g. colour, shape) and more abstract, or functional information (e.g. behaviour, use). The amount and salience of different feature types appears to differ systematically across categories (Cree and McRae, 2003, Warrington and McCarthy, 1983, Warrington and McCarthy, 1987, Warrington and Shallice, 1984). We propose that patterns of interactions among these variables also differ: for example, while concepts in the animal category have strong correlations among shared perceptual and "biological" functional properties of (e.g. between has eyes and can see, has a mouth and eats food), concepts in artifact categories such as tools have strong correlations among distinctive form (perceptual) and functional properties (Tyler *et al.*, 2000a, Tyler and Moss, 1997, 2001, see de Renzi and Lucchelli, 1994 for related accounts).

Thus, although there are differences between specific theories, the common theme of conceptual structure based accounts is that the combined effects of structural variables predict processing consequences that are picked up as psycholinguistic phenomena in lexical and semantic tasks, as well as the consequences of damage to the conceptual system in terms of the patterns of loss and preservation of different kinds of semantic information. Most recently, the approach has also been applied to the study of the neural basis of conceptual knowledge (Tyler *et al*., 2004, Tyler and Moss, 2001, Moss *et al*., 2005, Bright *et al*., 2005, Raposo *et al*., 2004, Taylor *et al*., in press). Although many other factors potentially reflect important aspects of conceptual representation (see Tranel *et al*., 1997 for discussion of many other candidates, including familiarity, age of acquisition and manipulability), the four variables referred to in the previous paragraph appear to be critically important to the internal structure of a concept: number of features, distinctiveness of those features, patterns of correlation among them and the interactions of these variables with feature type. In the following section we consider the theoretical development and empirical support for the role of each of these variables in turn.

3. CONCEPTUAL STRUCTURE VARIABLES: THEORETICAL DEVELOPMENT AND EMPIRICAL DATA

3.1. TYPE OF SEMANTIC FEATURE: PERCEPTUAL, FUNCTIONAL AND BEYOND

The distinction between perceptually grounded features and more abstract or "functional" features is a critical one in accounts of both the development (e.g. Madole *et al*., 1993, Mandler, 1992, Tyler *et al*., 2000a) and the final structure of the conceptual system (see Miller and Johnson-Laird, 1976, Barsalou, 1999 for discussions). Warrington and colleagues were amongst the first to explore distinctions among feature types in relation to the neural basis of semantic memory and the effects of damage to that system. They suggested that some categories, such as food and living things, are primarily differentiated in terms of their perceptual/sensory properties, while man-made objects are more reliant on their core functional properties. On the assumption that different types of features are stored in separate sub-systems within semantic memory, which is reflected in the topographical organisation of the neural substrate (Allport, 1985), Warrington and colleagues proposed that focal brain damage may disrupt one type of feature more than another, indirectly causing a greater impairment for any category of concept for which that feature type is particularly important (Warrington and McCarthy, 1983, Warrington and McCarthy, 1987, Warrington and Shallice, 1984, Borgo and Shallice, 1994, Farah *et al*., 1989, Farah and McClelland, 1991, Saffran, 2000).

These assumptions were instantiated in a connectionist model by Farah and McClelland (1991), although the difference between living things and non-living things was captured by the relative *amount* of functional and perceptual features for each concept, rather than their relative importance for differentiating members of the two domains, as had been proposed by Warrington and colleagues. Based on empirical data from a property identification study, the model's representations of non-living things contained similar numbers of functional and perceptual features, while living things had a far higher proportion of perceptual to functional features. When functional features were removed from the model, its performance was more impaired with artifacts than living things, while the reverse was true when perceptual features were removed from the model, demonstrating that category-specific deficit can emerge from a system organised by feature type rather than category *per se*.

Distinctions between feature types may have implications for on-line processing of conceptual information in the intact, as well as the impaired system; in a priming study, Moss *et al*. (1997a) found greater facilitation for targets denoting functional features of prime concepts (e.g. *tractor-farm*) than for targets denoting perceptual features (e.g. *tractor-wheels*), at least for artifact concepts such as tools and vehicles (primes referring to living things were not included in this study). This suggests more rapid activation of functional features, consistent with the notion that this type of feature is especially important for artifact concepts (see also Barton and Komatsu, 1989, Tyler and Moss, 1997). However, the interaction with other critical factors such as feature distinctiveness and correlation were not manipulated in this study. Moreover, although this sensory/functional or sensory/motor account of semantic memory has received considerable support from neuropsychological studies of patients with category-specific deficits for living things (e.g. Borgo and Shallice, 2001, de Renzi and Lucchelli, 1994, Basso *et al*., 1988, Sartori and Job, 1988, Moss *et al*., 1997b, see Saffran and Schwartz, 1994 for a review) and from neuro-imaging studies suggesting a topographical organisation of feature types in the brain (e.g. Martin and Chao, 2001, Martin *et al*., 2000, Hauk *et al*., 2004, Pecher *et al*., 2004, Pulvermüller, 1999), there have also been a number of important challenges to this approach. Many patients have been reported whose patterns of deficit appear inconsistent with the predicted associations between feature type and category, including those with a living things deficit who are impaired for both functional and perceptual properties (Caramazza and Shelton, 1998, Lambon Ralph *et al*., 1998, Moss *et al*., 1998), and those with deficits for visual properties who are not disproportionately impaired for living over non-living things (Lambon Ralph *et al*., 1998, see Capitani *et al*., 2003). Moreover, data from imaging studies concerning the proposed neural regions for different types of semantic feature are not consistent.

Finally, it is important to note that the broad distinction between perceptual and functional features is at best a convenient short-hand for the differentiation that may exist, based on multiple sensory, motor and associative sources of information. Finer-grained distinctions have been proposed, both in the original proposals of Warrington and colleagues and more recently (e.g. Crutch and Warrington, 2003, Cree and McRae, 2003, McRae and Cree, 2002, Tyler and Moss, 1997). For example, Barsalou *et al* (1999) cited in McRae and Cree (2002) developed a taxonomy of 28 "knowledge-types" which were broadly classified as 'entity', 'situation', 'introspective' and 'taxonomic'. Cree and McRae (2003), using a slightly modified version of this taxonomy, coded the knowledge type of features belonging to concepts in 34 categories spanning the living and nonliving domains. A principal components analysis revealed that eight of these knowledge types accounted for most of the variance in the distribution of feature types across categories: six 'entity' knowledge types ('entity behaviours', e.g., *clock - ticks*, 'external components', e.g. *tricycle – has pedals*, "made-of" features, e.g., *sink – made of enamel*, 'internal surface properties', e.g., *fridge – is cold*, 'external surface properties', e.g. *apple – is red*, and 'internal components', e.g., *cherry – has a pit*), and two 'situation' knowledge types ('functions', e.g., *tomato – eaten*, and locations, e.g. *cupboard – found in kitchens*, and ('entity' feature types). Moreover, the group of 'creature' categories could be distinguished from nonliving things based on the greater importance of entity behaviours to creatures and the greater importance of "function" and "made of" properties for nonliving things. However, while sensory features ('external components') were highly important to the representation of creatures, they did not distinguish between creatures and nonliving things, surprising considering the effects of sensory feature type shown in previous studies noted above. Perhaps more importantly, it is unclear how these knowledge types would pattern onto neural systems, as many refer to multicomponential cognitive processes (Cree and McRae, 2003). Nevertheless, the notion that there are different kinds of features which may contribute in different ways to the structure of concepts across various categories remains an important insight and variable within the conceptual structure framework, whether the assumption of neural organisation along this dimension is explicitly preserved (e.g. Devlin *et al.*, 1998) or not (e.g. Tyler *et al.*, 2000a). Moreover, the role of features that are grounded in the perceptual and motor systems has been critical in the development of recent "embodied theories" of conceptual representation (see Barsalou, 1999; also see Vinson & Vigliocco, this volume and Glenberg, this volume).

3.2. NUMBER OF FEATURES: FROM COMPLEXITY TO SEMANTIC RICHNESS

Perhaps the most basic claim of the conceptual structure framework is that some concepts contain more features than others. Theoretical linguistic analyses suggest, for example, that the concept *man* contains the features [+male +adult +human], while *bachelor* contains all these and [+unmarried]. A series of early memory and verification time studies failed to find any support for the *prima facie* prediction that words containing many features should be at a processing disadvantage compared to those with few features due to the additional complexity of their representations (e.g. Kintsch, 1974). In a classic study, Fodor, Garrett, Walker, & Parkes (1980) also failed to find evidence for decomposition, although they used a task in which subjects judged the relatedness of sentence elements as the index of complexity for simple (e.g. *to die*) vs complex verbs (e.g. *to kill* [i.e. to cause to die]), rather than a measure of on-line processing cost (see also Pitt, 1999 for a critique of the logic of this task) .The early results were quickly challenged on a number of grounds. Most importantly, it is not necessarily the case that a greater number of features should be problematic if we assume parallel rather than serial processing, and indeed, if the features have a high degree of connectivity among them, this may even facilitate processing (Gentner, 1981, and see discussion of correlations among features below).

More recently, it has become apparent that having more semantic features may indeed be beneficial to conceptual and lexical processing. Various empirical feature generation studies have demonstrated that concrete words typically have more features than do abstract words (e.g. de Mornay Davies and Funnell, 2000, Tyler *et al.*, 2002). This difference has been instantiated in a connectionist model designed to account for the effects of abstractness on the reading performance of patients with deep dyslexia (Plaut and Shallice, 1993). In this model, concrete words are represented over a greater number of microfeatures than are abstract words, and the activation of those features is also more consistent over different occurrences of the word. This leads to more stable patterns of activation, providing a basis for the better performance on concrete than abstract words when the input to the system is noisy, as in patients with deep dyslexia (Plaut and Shallice, 1993) or word meaning deafness (Tyler and Moss, 1997). Conversely, it is also possible that in some cases, concrete nouns may be impaired to a greater extent than abstract nouns and verbs – since the latter concepts contain fewer semantic features (especially perceptually based ones), they may be less adversely affected by damage to the semantic system (Breedin, Saffran & Costell, 1994; Saffran & Sholl, 1999).

The "richer" semantic representations for concrete words also provide an account of the processing advantage for these words in the normal system when coupled with a model of word recognition which allows activation from the semantic level to feed back to facilitate processing at lower levels (e.g. McClelland and Elman, 1986, Marslen-Wilson and Welsh, 1978, Balota *et al.*, 1991). It follows that the greater number of features activated for a given concept, the greater the amount of activation at the semantic level to facilitate semantic processing, and the more semantic activation is fed back to the orthographic and phonological levels to facilitate processing here (Tyler *et al.*, 2000b). These claims have recently been tested in greater detail in a series of studies by Pexman and colleagues (Pexman *et al.*, 2002, 2003). For example, Pexman *et al* (2003) showed that semantic (concrete/abstract) decisions to words with a greater number of features (NoF) were faster than to low NoF words, a finding attributed to the greater semantic activation of high NoF words. The purported facilitatory effects of semantic feedback activation were confirmed in naming and lexical decision tasks, in which reaction times were faster to high than to low NoF words (Pexman *et al.*, 2002), and to polysemous than nonpolysemous words (Hino and Lupker, 1996). However, semantic feedback activation is not always advantageous: words with many synonyms purportedly result in a large amount of semantic activation which feeds back to several different orthographic representations, creating *competition* at this level. Indeed, lexical decisions to words with many synonyms are slower than to words without synonyms (Pecher, 2001). Taken together, these elegant series of studies provide compelling evidence for a distributed, componential semantic system in which NoF plays an important role. However, it should be noted that most of the evidence on this point comes from the study of noun concepts, and similar results have yet to be demonstrated for other form classes, including verbs. In fact, one recent study suggests that (contra Fodor et al, 1980) there may indeed be a processing cost for complex event verbs over simple state verbs in reading and lexical decision times (Gennari & Poeppel, 2003).

Although semantic richness or number of features is clearly an important aspect of conceptual structure, it can only be part of the story. As discussed above, concrete words typically have a greater number of features than do abstract words, affording several processing advantages in the healthy system, and providing greater feedback compensation when lexical systems are damaged. However, within the domain of concrete words, living things typically have more features than do artifacts, as shown in property generation norms (Randall *et al.*, 2004, Garrard *et al.*, 2001, Greer *et al.*, 2001). Whether living things show a consistent processing advantage over non-living things in the intact system is debatable (Pilgrim *et al.*, 2005, Laws and Neve, 1999, Gaffan and Heywood, 1993). The pattern varies with task demands; for example, in studies using pictures, living things are disadvantaged when fine-grained distinctions among similar items are needed, but this effect can be removed or even reversed when colour and texture information is added to the picture – features that are highly informative for living things (Price and Humphreys, 1989, Moss *et al.*, 2005). Similarly, level of categorisation is important, with an advantage for living things often revealed for category level identification, but a disadvantage for basic level naming (Humphreys *et al.*, 1988, Lloyd-Jones and Humphreys, 1997, see Moss *et al.*, 2005 for a discussion of potential neural bases of these differences). However, for brain damaged patients, there is commonly a clear disadvantage for living things, which remains even when familiarity and other potentially confounding factors are taken into account. Thus, a high number of features alone does *not* protect these concepts. The basis for this discrepancy between the concrete/abstract dissociation and the living/artifact dissociation may lie in the contrasting loci of damage in these patients (both functionally and neurally). Many patients with deficits for abstract words have impairments affecting lexical systems, often limited to a specific modality of input or output (e.g. deep dyslexia, word meaning deafness). Their patterns of performance over semantic domains can be accounted for as the relative success of feedback from a largely intact semantic system to an impaired lexical system; concepts with many features providing greater support. Patients with living things deficits, on the other hand, typically have impairments to central conceptual systems, not limited to a specific modality of input or output. Although living things concepts have many features, it is the nature of those features, both in terms of their distinctiveness (or lack thereof) and their correlations, that make them vulnerable to damage, as discussed in the following sections (see also Vigloccio & Vinson, this volume for a related discussion of differences in the representation of abstract and concrete words).

3.3. FEATURE DISTINCTIVENESS: FROM CUE VALIDITY TO RELEVANCE

Feature distinctiveness essentially refers to the number of concepts in which a feature appears, ranging from one (e.g. *has an udder*, which is true of cows) to very many (e.g. *has eyes*, true of all animals). This factor can be traced back to the notion of cue validity – the conditional probability with which a feature signals a specific concept (Rosch and Mervis, 1975). Similarly, Devlin *et al* (1998) characterise features in terms of their informativeness – features occurring for very few concepts are highly informative in identifying specific concepts, while those occurring for many concepts place few constraints on the range of possible concepts to which the feature belongs.

Although distinctiveness can be readily captured in distributed connectionist models of conceptual structure in terms of the overlap of activation of feature units within concept representations (e.g. Devlin *et al*., 1998, Durrant-Peatfield *et al*., 1997, Tyler *et al*., 2000a, Greer *et al*., 2001), distinctiveness has also been an important issue within the very different framework of localist hierarchical models of semantic memory. Here the drive for cognitive economy suggested that features shared by all members of a category be stored at higher levels of the hierarchy only, being "inherited" by concepts within that category, rather than duplicated for each member. Distinctive features on the other hand would need to be represented individually for each feature at lower levels (Collins and Loftus, 1975, see Moss and Marslen-Wilson, 1993 for a discussion). This was the framework in which Warrington (1975) initially interpreted the finding that patients with semantic deficits have more difficulty with distinctive than general properties, claiming that access to semantic memory may proceed in a "top down" manner starting with the general information stored at the category level, culminating in the most distinctive properties at the ends of the branches – these later, more detailed retrieval processes being more susceptible to disruption.

Although the weight of evidence does not support the claim of general to specific access (Rapp and Caramazza, 1989) or the hierarchical model more generally (see introduction), the key finding that distinctive properties are more vulnerable to damage than are shared properties has proved to be a defining characteristic of semantic impairments (e.g. Bub *et al*., 1988, Hodges *et al*., 1995, Moss *et al*., 1997a, Moss *et al*., 1998, Hart and Gordon, 1992, Tippett *et al*., 1995). This general pattern can be accounted for within a distributed connectionist framework, given that distinctive features are experienced less frequently than are highly shared properties (which occur with many different objects) resulting in weaker connection strengths over all.

However, effects of distinctiveness are unlikely to be uniform across the conceptual system. Firstly, there is considerable evidence that concepts differ in terms of the distinctiveness of their features; specifically, living things appear to have a higher proportion of shared to distinctive features than do artifact concepts. For example, Randall *et al* (2004) analysed the distribution of the distinctiveness of properties generated by a group of participants to 93 concepts from the categories of animals, fruit, tools and vehicles. We measured the distinctiveness of a feature as an inverse function of the number of concepts for which it is generated. Each feature has a distinctiveness value associated with it ranging from 1 (highly distinctive) to 0 (not distinctive). As predicted, the mean distinctiveness of features within artifact concepts was significantly greater than that for living things. This effect has been found in a number of property generation studies, although the cut-off between distinctive and shared properties is defined in various ways across studies (Devlin *et al*., 1998, Garrard *et al*., 2001, McRae and Cree, 2002, Vinson and Vigliocco, 2002) and although the precise proportions vary across studies, largely due to the tendency for participants to under-produce highly shared features unless encouraged to do so (see Rogers *et al*., 2004, for a comparison across studies which highlights the potential for variation in resulting feature sets). While this domain difference in shared/distinctive feature ratio has important consequences for the effect of damage on the conceptual system, the precise consequences of the ratio of distinctive to shared properties across domains cannot be considered in isolation from other factors with which it interacts, most importantly feature type and feature correlation.

Recently, Sartori & Lombardi (2004) have proposed a new feature variable of semantic relevance. Although very similar to the notion of distinctiveness, this variable is weighted according to the importance of the feature for the meaning of the concept (as derived from the number of responses in a property generation listing study and calculated using a relevance matrix). Thus, while distinctiveness is a concept-independent measure (the distinctiveness of a feature will be the same for all those concepts in which it occurs), relevance is concept-dependent (the same feature may be more relevant for one concept than another – for example, in Sartori & Lombardi's dataset, the relevance of *has a beak* is greater for the concept *duck* than *swan*, largely reflecting the fact that more participants listed this feature for duck in the generation study). Sartori & Job's analyses of the distribution of feature relevance values across categories suggest a similar pattern to that found for distinctiveness in the earlier studies; most importantly, features of living things (and especially fruit and vegetables) had significantly lower relevance values than those of non-living things. The relevance measure may be an important development on the notion of distinctiveness as it captures the relationship between a concept and its features as a graded one, rather than in an all-or-none manner (see also Vinson and Vigliocco, 2002, for a related approach to weighting features by salience).

3.4. Feature Correlation: Clusters and mutual activation

Rosch *et al* (1976) observed that properties of natural categories, rather than being independent, tend to cluster together, for example, that creatures with feathers generally have wings, beaks and lay eggs. Certain combinations of properties occur together much more frequently than do others. In a series of studies, Keil (1986) demonstrated that clusters of properties are larger and more densely intercorrelated for concepts within the domain of living things than of manmade objects. This notion of correlation of properties, and their variation across domains of concepts, is a central tenet of the conceptual structure approach.

However, early studies of the role of property correlation in real world categories (as opposed to learning artificial concepts) were rather mixed. Malt & Smith (1984) generated a property correlation matrix for a set of basic level concepts, based on participants' responses in feature generation and verification studies. They found that the incidence of property correlations was much greater than would be expected by chance, with about a third of all potential pairs correlated at greater than the .05 level. However, it was less clear that participants actually used this information to perform concept processing tasks: for example, adding property correlation information to a simple family resemblance weighted sum model significantly improved predictions of ratings in a typicality rating task only under certain conditions - when highly salient correlations were considered and explicit comparisons between correlated and uncorrelated pairs were required. This finding, along with other empirical results showing that correlations among features had little or no effects on explicit off-line category learning tasks (Murphy and Wisniewski, 1989), led to claims that conceptual processing is not typically sensitive to statistical regularities such as feature correlation, but rather that correlations tend only to be noticed when they are explicitly pointed out, or when participants are aware of a theoretical basis for why certain properties might co-occur (Murphy and Medin, 1985).

This rather negative view of the role of correlation has recently been challenged by a number of studies demonstrating the importance of intercorrelation and distinctiveness in conceptual structure, both in predicting patterns of semantic impairment following brain damage and on-line activation of information in the intact system. In each case, the theoretical proposals have also been implemented in distributed connectionist models to test the validity of the major assumptions (McRae *et al*., 1997, McRae *et al*., 1999, McRae and Cree, 2002, Devlin *et al*., 1998, Durrant-Peatfield *et al*., 1997, Moss *et al*., 1998, Tyler *et al*., 2000a, Tyler and Moss, 2001).

McRae & colleagues (McRae *et al*., 1997, McRae *et al*., 1999) suggest that correlations among semantic features play a role in the early computation of word meaning in on-line tasks – in contrast to the more metalinguistic conceptual reasoning tasks that had been used in earlier studies, in which higher level theoretical knowledge may be more relevant (e.g. Keil, 1989, Rips, 1989). To investigate this issue, McRae *et al* compared the impact of featural variables in fast on-line tasks (e.g. semantic priming, speeded feature verification) with that in slower, untimed tasks, more akin to those used in earlier studies (e.g. similarity and typicality ratings). Featural variables for concepts were established in a large-scale property generation study. Analysis of these norms supported the claim that living things have a significantly greater number of correlated properties than do artifacts (Keil, 1989), although the overall number of features did not differ across domains in this cohort.

Results from the on-line semantic tasks suggested that the initial computation of word meaning is indeed highly sensitive to the distributional statistics of features within concepts. For example, in a short SOA priming task, facilitation increased as a function of the overlap in individual features for artifact pairs (i.e., number of shared features; e.g. *pistol-rifle*), while overlap specifically in correlated features predicted facilitation for living things (e.g. *eagle-hawk*). However, in an untimed similarity rating task, the effect of correlation for living things disappeared, suggesting that this factor affects initial activation of the meaning rather than the eventual stable state. A similar pattern was shown in a feature verification paradigm: participants were asked to indicate whether a feature was true of a concept (e.g. *deer-hunted by people*). Feature correlation was manipulated such that half of the features were highly correlated with other features of the concept (e.g. *hunted by people* is correlated with many other properties of *deer)* while for the other half of the stimuli the features were presented with concepts where they were weakly correlated with other features (e.g. *duck –hunted by people*). Feature correlation was a significant predictor of reaction times, over and above other important factors such as conceptual familiarity and production frequency of the feature. This finding was replicated in McRae *et al* (1999), who also found a significant, albeit smaller effect of correlation strength at a longer SOA, and by Randall *et al* (2004), who reported significantly slower reaction times to the weakly correlated, distinctive features of living things than to the shared properties of living things and the relatively strongly intercorrelated properties of non-living things in a speeded feature verification task. However, correlation strength was not a significant predictor of typicality ratings in an untimed task

(McRae *et al*., 1999) nor an unspeeded feature verification task (Randall *et al*., 2004). Finally, McRae *et al* simulated the main effects of these two empirical studies in a distributed connectionist model which mapped from word form units to semantic representations which were distributed over feature units, directly reflecting the structure of concepts derived from the property generation study. In a simulation of the priming study, the model replicated the behavioural results, showing an early effect of overlap of correlated features for living things, but of individual feature similarity for artifacts. These findings suggest that one reason for null effects of correlation in the earlier category learning studies may have been due at least in part to their extended time course, which would not pick up the early effects on meaning activation that were so clearly revealed in priming and speeded feature verification tasks.

4. CONCEPTUAL STRUCTURE ACCOUNT REVISITED: CORRELATION, DISTINCTIVENESS, FEATURE TYPE AND DOMAIN

Our CSA model described earlier in this chapter also stresses the combined contribution of feature correlations and distinctiveness in determining conceptual structure. A critical difference between our approach and that of Gonnerman and colleagues (Gonnerman *et al*., 1997, Devlin *et al*., 1998) and McRae and colleagues (McRae *et al*., 1997, Cree and McRae, 2003) is that we incorporate a set of claims about how these variables interact with differences in feature type - specifically form (perceptual properties) and function - in the living and non-living domains. These claims draw on the developmental literature, which investigates how children learn the relations among properties of concepts. We claim that an essential aspect of conceptual structure is the pattern of correlations between form and function (Tversky and Hemenway, 1984). If a perceptual form is consistently observed performing a function, then a system which is sensitive to co-occurrences will learn that a specific form implies a specific function (Madole *et al*., 1993, Mandler, 1992). The nature of these form-function relations distinguishes between living things and artifacts. Artifacts have distinctive forms, which are consistently associated with the functions for which they were created (de Renzi and Lucchelli, 1994, Keil, 1986, 1989, see also Caramazza *et al*.'s (1990) claim of privileged relations among properties for a similar view). Artifacts are generally designed to perform a single distinctive function so that their form is as distinctive as the function. In contrast, living things tend to 'do' similar things and they tend to resemble each other, thus they share many features. Individual variations in form tend not to be functionally significant (e.g. *a lion's mane*). Even so, living things (like artifacts) also have form-function correlations. But whereas the form-function correlations for artifacts involve distinctive properties, for living things it is the shared properties (e.g. *eyes, legs*) that are involved in form-function correlations (e.g. *eyes-see; legs-move*). We refer to these as biological functions (Durrant-Peatfield *et al*., 1997, Tyler *et al*., 2000a, Tyler and Moss, 1997). Unlike the sensory/functional account, we do not claim that functional information is more important for artifacts than for living things, but rather that there is a difference across domains in the *kind* of functional information that is most strongly correlated – and therefore most robust to damage. Living things have many, very important functional properties, but the most important ones concern their biological activities that are frequently shared across most or all members of a category, rather than their intended use or purpose in relation to human beings (Tyler and Moss, 1997). These predictions were supported by the analysis of property generation norms (Randall *et al*., 2004). First, living things concepts had more features that were significantly more correlated with each other than were concepts in the non-living domain, a finding that has also been reported for several other property norm studies (McRae *et al*., 1997, Garrard *et al*., 2001, Vinson *et al*., 2003, Devlin *et al*., 1998). Most importantly for the CSA account, the distinctiveness of features that participated in form-function correlations (e.g., has *a blade – is used for cutting)* was significantly greater for concepts in the non-living than the living domain. However, this finding has not been replicated in other property norm studies. Garrard *et al* (2001) and Vinson *et al* (2003) both reported a greater number of correlated features for living than non-living things, *but* that distinctive properties of living things were more, rather than less correlated than those of non-living things. While it is possible that both these studies were limited by the small number of concepts entered into the analyses, and by the under-estimation of shared properties overall, it will clearly be important to establish in future studies whether the interaction between distinctiveness, correlation and domain is a robust one, as claimed by the CSA.

4.1. The Conceptual Structure Account and Semantic Deficits

Taken together, the assumptions of the CSA predict that patients with category specific deficits for living things will show a particular pattern of loss and preservation of features; specifically, that they will have the greatest problem with the distinctive properties of living things, due to the inherently weak correlations of this type of information. Since most semantic tasks require within-category discriminations that rely on intact knowledge of distinctive properties, this will generally show up as a deficit for living things, even though knowledge of the highly correlated shared information may be intact – indeed in some cases may be superior to knowledge of shared information about man-made objects. This pattern was demonstrated in a detailed single-case study of a patient with a deficit for living things following Herpes Simplex Encephalitis (HSE). Across a range of tasks, including property verification, sorting and naming to definition, RC showed selective difficulties with the distinctive properties of living things, but no apparent difficulty for shared properties. We have reported similar findings for other patients with living things deficits with an aetiology of HSE (Tyler and Moss, 2001, Moss *et al*., 2002). Other authors have also reported relatively preserved knowledge of shared properties and impaired knowledge of distinctive properties of living things for patients with category-specific deficits (e.g. patient EW; Caramazza and Shelton, 1998). For most other patients in the literature, the appropriate contrasts between distinctive and shared information were not tested, so it is not possible to determine whether these patients were more impaired with the distinctive properties of living things. However, in several reports, there are hints that this is the case (e.g. Sartori and Job, 1988). Similarly, Sartori & Lombardi (2004) report evidence which suggests that the typically low semantic relevance of features of living things cause patients' difficulties with concepts in this domain: when relevance was controlled across domain, living thing deficits were reduced or even reversed.

4.2. The Conceptual Structure Account and Speed of Intact Processing

The CSA also makes predictions for the speed with which different types of features become activated during normal on-line processing in the intact system. This includes the prediction that there will be a disadvantage in rapidly activating the distinctive properties of living things relative to other kinds of features, due to their lack of correlation with other information – even in the normal conceptual system - on the premise that mutual activation produces faster initial processing times for correlated features. For shared properties we predicted little difference across domains, since both living and non-living things have groups of correlated shared properties – if anything, the pattern should be in the opposite direction, as living things have a greater number of shared properties. These predictions were supported by the results of a speeded feature verification task, which was designed to tap into the early stages of semantic activation, with short presentation times, backward masking and a response deadline at only 450 msec SOA. There was a significant interaction between distinctiveness and domain; feature verification for distinctive features of living things was disproportionately slow and error-prone, while there was no difference between living and non-living things for shared features (Randall *et al.*, 2004). These results were simulated in a distributed connectionist model, which mapped orthographic word form units, via a set of hidden units, onto semantic representations which were distributed over semantic feature units, directly reflecting the structure of the concepts derived from our property generation study. The model replicated the domain by distinctiveness interaction in the speeded feature verification results, showing that the distinctive properties of living things were more error prone and took longer to settle than the distinctive properties of non-living things, while there were no domain differences for shared properties. These results demonstrated that processing differences across domains arise on the basis of differences in the correlational structure of concepts within domains.

5. CHALLENGES FOR CONCEPTUAL STRUCTURE FRAMEWORK

Conceptual structure accounts have taken us a long way in understanding the psychology and neuropsychology of conceptual representation and processing. However, as with any model of conceptual representation, this account faces several challenges. For example, the CSA makes clear predictions about the relationship between the severity of impairment and the degree of deficit for living and non-living things. Based on the results of our computational simulations, we predicted that at most levels of damage, living things would be most impaired, but when damage was very severe, and overall performance very inaccurate, that living things would have a slight processing advantage over artifacts (Tyler *et al.*, 2000a, Greer *et al.*, 2001). We have argued that this pattern is due to the large number of highly intercorrelated shared properties for living things, which are the only information that can withstand a high degree of damage, allowing a small percentage of living things trials to be correct (see Moss and Tyler, 2000 for further details). Our own longitudinal studies of patients with degenerative diseases affecting the semantic system have supported this hypothesis (Moss and Tyler, 2000, Moss *et al.*, 2002). Moreover, some of the handful of patients reported with artifact deficits do seem to have very severe deficits, consistent with this view (e.g. VER and YOT, Warrington and McCarthy, 1983, 1987). However, data from cross-sectional studies of groups of AD patients are mixed, with some showing no relationship between severity and domain effects (Garrard *et al.*, 1998). Other patients with category-specific semantic impairments for non-living things have been reported that were not severely impaired in other domains (Hillis and Caramazza, 1991, Sacchett and Humphreys, 1992). The domain by severity interaction thus remains controversial (Garrard *et al.*, 1998, Caramazza and Mahon, 2003).

A further challenge to distributed conceptual structure accounts such as the CSA is to reconcile their claim of a unitary semantic system with the behavioural and functional neuro-imaging evidence for a sensory-motor organisation of features (e.g. Warrington and Shallice, 1984, Chao *et al.*, 1999, Martin *et al.*, 2000). Some authors have addressed this issue by proposing a hybrid model of conceptual structure which includes a sensory-motor distribution of semantic features (e.g. the Featural and Unitary Semantic Space (FUSS) model, Vigliocco *et al.*, 2004, and Cree and McRae, 2003, where knowledge types play a critical role alongside conceptual structure variables). According to the FUSS model, semantic representations of object (and action) knowledge are represented at two levels: a conceptual feature space organised by sensory-motor feature type, and a second, “lexico-semantic” space organised by the conceptual structure factors of feature salience, sharedness/distinctiveness, and intercorrelation, which binds features from the conceptual feature space. This architecture allows for isolated impairments of specific kind of sensorimotor feature and thus category for which these features are particularly important, as well as patterns of deficits consistent with a conceptual structure approach, by hypothesising damage to the conceptual feature and lexico-semantic space, respectively. Further studies will need to determine the extent to which the FUSS's two-tiered model of conceptual representation is supported by functional-neuroanatomical activation patterns.

Simmons and Barsalou (2003) likewise argued that a more comprehensive and powerful model of semantic memory may be achieved by integrating conceptual structure and sensory-functional or sensory-motor accounts, and in this vein developed their Conceptual Topography Theory (CTT). The CTT draws heavily on the non-human primate literature on object processing, as well as Damasio's (1989) convergence zone theory of human object processing. It postulates that each sensory and motor system contains "feature maps" processing the respective elementary object features. These sensory and motor features are bound together into increasing more complex feature conjunctions from posterior to anterior regions in a hierarchical system of "convergence zones" (association areas) in each sensory and motor stream. The anteromedial temporal lobe is proposed to play a special role in object processing in that it purportedly processes the most complex conjunction of visual features as well as the multimodal feature conjunctions. The CTT makes two additional assumptions about the neural code of conceptual similarity. Specifically, it postulates that the proximity of neurons in a convergence zone increases as a function of the similarity of the features they conjoin (the "similarity-in-topography principle"), and that clumps of conjunctive neurons representing category members become more dispersed as the similarity of the represented category members decreases (the "variable dispersion principle"). Thus, the CTT's sensory and motor feature maps correspond to these feature representations in the sensory-functional and sensory-motor accounts (Warrington and McCarthy, 1983, Warrington and McCarthy, 1987, Martin and Chao, 2001, and Martin *et al.*, 2000 respectively). The CTT also instantiates distinctiveness in neural space: concepts which share many features (e.g. living things) would be represented close together in the convergence zone (the "similarity-in-topography principle"), while the representations of concepts which share fewer features (e.g. nonliving things) would be more dispersed in the convergence zone (the "variable dispersion principle").

We recently tested the central claims of this "neurocognitive" account in a series of functional neuro-imaging studies (Tyler *et al.*, 2004, Moss *et al.*, 2005). Healthy participants were instructed to perform two different picture-naming tasks with the same picture stimuli. In a domain-level naming task, participants silently named the domain (i.e., "living", "manmade") to which pictured objects belonged, while a basic-level naming task with the same picture stimuli required participants to silently name the pictured object (e.g., "tiger"). We hypothesised that the domain-level naming task would require relatively simple visual feature conjunctions (e.g., curvature) mediated by more posterior regions in the visual object processing stream to differentiate living from nonliving things, while the basic-level naming task would require relatively more complex visual feature conjunctions mediated by more anterior regions in the visual object processing stream in order to distinguish the pictured object from other, visually similar objects (e.g., a tiger from a lion). Consistent with these hypotheses, domain-level naming was associated with more posterior ventral occipitotemporal, and basic-level naming additionally with more anteromedial temporal lobe activity (Tyler *et al.*, 2004). We next hypothesised that since living things are characterised by many shared and relatively few distinctive features while nonliving things typically have a greater proportion of distinctive to shared features, that the identification of living things would be more visually demanding, requiring more complex visual feature conjunctions compared to the basic-level naming of nonliving things. As predicted, we found greater anteromedial temporal lobe activity presumably reflecting complex visual feature conjunctions during the basic-level naming of living compared to matched sets of nonliving things, a finding confirmed in behavioural studies with patients with brain damage including the anteromedial temporal lobe (Moss *et al*., 2005).

The ultimate challenge to conceptual structure accounts may be to explain how it is instantiated in the brain. Neurocognitive accounts, which integrate psycholinguistically and neuropsychologically verified principles of conceptual structure with a hierarchical object processing model developed in non-human primates, may have the capacity to meet this challenge. We hope that future psycholinguistic, neuropsychological and neuro-imaging research on conceptual structure will continue to be conducted in parallel, mutually enriching each other's findings, to determine how feature types, NoF, distinctiveness, intercorrelation and the combination of these factors are behaviourally and neurally represented and processed.

ACKNOWLEDGEMENTS

This work was supported by an MRC programme grant to LKT and a grant from the Roche Research Foundation and Olga Mayenfisch Foundation to KIT.

References

Allport DA (1985). Distributed memory, modular subsystems and dysphasia. In Newman, SK & Epstein, R (Eds.) *Current perspectives in dysphasia.* Edinburgh, Churchill Livingstone.

Balota DA, Ferraro RF and Connor LT (1991). On the early influence of meaning in word recognition: A review of the literature. In Schwanenflugel, PJ (Ed.) *The psychology of word meanings.* Hillsdale, NJ, Erlbaum.

Barsalou LW (1999). Perceptual symbol systems. *Behavioral and Brain Sciences,* 22**,** 577-609.

Barsalou LW, Solomon KO and Wu LL (1999). Perceptual simulation in conceptual tasks. In Hiraga, MK, Sinha, C & Wilcox, S (Eds.) *Cultural, typological, and psychological perspectives in cognitive linguistics: The proceedings of the 4th conference of the International Cognitive Linguistics Association.* Amsterdam, John Benjamins.

Barton ME and Komatsu L (1989). Defining features of natural kinds and artifacts. *Journal of Psycholinguistic Research,* 18**,** 433-447.

Basso A, Capitani E and 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.

Borgo F and Shallice T (1994). When living things and other 'sensory quality' categories behave in the same fashion: a novel category specificity effect. *Neurocase,* 7**,** 201-220.

Borgo F and Shallice T (2001). When living things and other sensory quality categories go together: a novel category-specific effect. *Neurocase,* 7**,** 201-220.

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

Bright P, Moss HE, Stamatakis EA and Tyler LK (2005). The anatomy of object processing: The role of anteromedial temporal cortex. *Quarterly Journal of Experimental Psychology, Section B,* 58**,** 361-377.

Bub DN, Black SE, Hampson E and Kertesz A (1988). Semantic encoding of pictures and words: some neuropsychological observations. *Cognitive Neuropsychology,* 5**,** 27-66.

Bunn EM, Tyler LK and Moss HE (1998). Category-specific semantic deficits: The role of familiarity and property type reexamined. *Neuropsychology,* 12**,** 367-379.

Capitani E, Laiacona, , Mahon, B and Caramazz, A (2003). What are the facts of semantic category-specific deficits? A critical review of the clinical evidence. *Cognitive Neuropsychology,* 20**,** 213-261.

Caramazza A and Hillis AE (1991). Lexical organisation of nouns and verbs in the brain. *Nature,* 349**,** 788-790.

Caramazza A, Hillis AE, Rapp BC and Romani C (1990). The multiple semantics hypothesis: Multiple confusions? *Cognitive Neuropsychology,* 7**,** 161-189.

Caramazza A and Mahon BZ (2003). The organization of conceptual knowledge: the evidence from category-specific semantic deficits. *Trends in Cognitive Sciences,* 7**,** 354-361.

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

Chao LL, Haxby JV and Martin A (1999). Attribute-based neural substrates in temporal cortex for perceiving and knowing about objects. *Nature Neuroscience,* 2**,** 913-919.

Collins AM and Loftus EF (1975). A spreading activation theory of semantic processing. *Psychological Review,* 82**,** 407-428.

Cree, GS and McRae K (2003). Analyzing the factors underlying the structure andcomputation of the meaning of chipmunk, cherry, chisel, cheese, and cello (and many other such concrete nouns). *Journal of Experimental Psychology: General,* 132**,** 163-201.

Cree GS, McRae K and McNorgan C (1999). An attractor model of lexical conceptual processing: simulating semantic priming. *Cognitive Science,* 23**,** 371-414.

Crutch SJ and Warrington EK (2003). The selective impairment of fruit and vegetable knowledge: a multiple processing channels account of fine-grain category specificity. *Cognitive Neuropsychology,* 20**,** 355-372.

Damasio AR (1989). Time-locked multi-regional retro-activation: A systems-level proposal for the neural substrates of recall and recognition. *Cognition,* 33**,** 25-62.

de Almeida RG (1999). What do category-specific semantic deficits tell us about the representation of lexical concepts? *Brain and Language,* 68**,** 241-248.

de Mornay Davies P and Funnell E (2000). Semantic representation and ease of predication. *Brain and Language,* 73**,** 92-119.

de Renz, E and Lucchelli F (1994). Are semantic systems separately represented in the brain? The case of living category impairment. *Cortex,* 30**,** 3-25.

Devlin JT, Gonnerman LM, Andersen ES and Seidenberg MS (1998). Category-specific semantic deficits in focal and widespread brain damage: A computational account. *Journal of Cognitive Neuroscience,* 10**,** 77-94.

Dixon M, Bub DN and Arguin M (1997). The interaction of object form and object meaning in the identification performance of a patient with category-specific visual agnosia. *Cognitive Neuropsychology,* 14**,** 1085-1130.

Durrant-Peatfield MR, Tyler LK, Moss HE and Levy JP (1997). The distinctiveness of form and function in category structure: A connectionist model. In Shafto, MG & Langley, P (Eds.) *Proceedings of the19th Annual Conference of the Cognitive Science Society.* Mahwah, NJ, Lawrence Erlbaum Associates.

Farah MJ, Hammond KM, Mehta Z and Radcliff G (1989). Category-specificity and modality-specificity in semantic memory. *Neuropsychologia,* 27**,** 193-200.

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

Fodor JA and Fodor JD (1980). Functional structure, quantifiers and meaning postulates. *Linguistic Inquiry,* 11**,** 4.

Fodor JA, Garrett MF, Walker ECT and Parkes CH (1980). Against definitions. *Cognition*, 8, 263–367.

Forde EME, Francis D, Riddoch MJ, Rumiati RI and Humphreys GW (1997). On the links between visual knowledge and naming: A single case study of a patient with a category-specific impairment for living things. *Cognitive Neuropsychology,* 14**,** 403-458.

Forde EME and Humphreys GW (1999). Category-specific recognition impairments: a review of important case studies and influential theories. *Aphasiology,* 13**,** 169-193.

Franklin S (1989). Dissociations in auditory word comprehension: evidence from 9 'fluent' aphasic patients. *Aphasiology,* 3**,** 189-207.

Gaffan D and Heywood CA (1993). A spurious category-specific visual agnosia for living things in normal human and non-human primates. *Journal of Cognitive Neuroscience,* 5**,** 118-128.

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

Garrard P, Patterson K, Watson PC and Hodges JR (1998). Category-specific semantic loss in dementia of the Alzheimer's type: Functional-anatomical correlations from cross-sectional analyses. *Brain,* 121**,** 633-646.

Gennari S and Poeppel D (2003).Processing correlates of lexical semantic complexity. *Cognition*, 89, 27-41.

Gentner D (1981). Some interesting differences between nouns and verbs. *Cognition and Brain Theory,* 4**,** 161-178.

Gonnerman LM, Andersen ES, Devlin JT, Kempler D and Seidenberg MS (1997). Double dissociation of semantic categories in Alzheimer's disease. *Brain and Language,* 57**,** 254-279.

Greer MJ, van Casteren M, McLellan, J *et al.* (2001). The emergence of semantic categories from distributed featural representations. In Moore, JD & Stenning, K (Eds.) *Proceedings of the 23rd Annual Conference of the Cognitive Science Society.* London, UK, Lawrence Erlbaum Associates.

Hart J and Gordon B (1992). Neural subsystems for object knowledge. *Nature,* 359**,** 60-64.

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

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

Hino Y and Lupker SJ (1996). Effects of polysemy in lexical decision and naming: An alternative to lexical access accounts. *Journal of Experimental Psychology: Human Perception and Performance,* 22**,** 1331-1356.

Hinton GE, McClelland JL and Rumelhart DE (1986). Distributed representations. In McClelland, JL & Rumelhart, DE (Eds.) *Parallel distributed processing: Explorations in the microstructure of cognition.* Cambridge, MA, MIT Press.

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

Hodges JR, Graham N and Patterson K (1995). Charting the progression in semantic dementia: Implications for the organisation of semantic memory. *Memory,* 3**,** 463-495.

Humphreys GW, Riddoch MJ and Quinlan P (1988). Cascade processes in picture identification. *Cognitive Neuropsychology,* 5**,** 67-103.

Jackendoff R (1983). *Semantics and cognition,* Cambridge, MA, MIT Press.

Jackendoff R (1989). What is a concept, that a person can grasp it? *Mind and Language,* 4**,** 68-102.

Kawamoto AH (1993). Nonlinear dynamics in the resolution of lexical ambiguity: a parallel distributed processing account. *Journal of Memory and Language,* 32**,** 474-516.

Keil FC (1986). The acquisition of natural kind and artifact terms. In Domopoulous, W & Marras, A (Eds.) *Language learning and concept acquisition.* Norwood, NJ, Ablex.

Keil FC (1989). *Concepts, kinds and cognitive development,* Cambridge, MA, MIT Press.

Kintsch W (1974). *The representation of meaning in memory,* Hillsdale, Lawrence Erlbaum.

Komatsu LK (1992). Recent views on conceptual structure. *Psychological Bulletin,* 112**,** 500-526.

Laiacona M and Caramazza A (2004). The noun/verb dissociation in language production: Varieties of causes. *Cognitive Neuropsychology,* 21**,** 103-123.

Lambon Ralph MA, Howard D, Nightingale G and Ellis AW (1998). Are living and non-living category-specific deficits causally linked to impaired perceptual or associative knowledge? Evidence from a category-specific double dissociation. *Neurocase,* 4**,** 311-338.

Laws KR and Neve C (1999). A 'normal' category-specific advantage for naming living things. *Neuropsychologia,* 37**,** 1263-1269.

Lloyd-Jones T and Humphreys GW (1997). Categorizing chairs and naming pears: Category differences in object processing as a function of task and priming. *Memory and Cognition,* 25**,** 606-624.

Madole KL, Oakes LM and Cohen LB (1993). Developmental changes in infants' attention to function and form-function correlations. *Cognitive Development,* 8, 189-209.

Malt BC and Smith E (1984). Correlated properties in natural categories. *Journal of Verbal Learning and Verbal Behavior,* 23, 250-269.

Mandler JM (1992). How to build a baby II: Conceptual primitives. *Psychological Review,* 99, 587-604.

Marslen-Wilson WD and Welsh A (1978). Processing interactions and lexical access during word-recognition in continous speech. *Cognitive Psychology,* 10, 29-63.

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

Martin A, Ungerleider LG and Haxby JV (2000). Category-specificity and the brain: The sensory-motor model of semantic representations of objects. In Gazzaniga, MS (Ed.) *The cognitive neurosciences.* 2nd ed. Cambridge, MIT Press.

Masson M (1995). A distributed memory model of semantic priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition,* 21, 3-23.

McClelland J and Elman J (1986). The TRACE model of speech perception. *Cognitive Psychology,* 18, 1-86.

McNamara TP (1992). Priming and constraints it places on theories of memory and retrieval. *Psychological Review,* 99, 650-662.

McRae K and Cree GS (2002). Factors underlying category-specific semantic deficits. In Forde, EME & Humphreys, GW (Eds.) *Category-specificity in brain and mind.* Hove, UK, Psychology Press.

McRae K, Cree GS Westmacott, R and de Sa, VR (1999). Further evidence for feature correlations in semantic memory. *Canadian Journal of Experimental Psychology,* 53, 360-373.

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

Miller G and Johnson-Laird PN (1976). *Language and perception,* Cambridge, MA, Harvard University Press.

Miller GA (1978). Semantic relations among words. In Halle, M, Bresnan, J & Miller, GA (Eds.) *Linguistic theory and psychological reality.* Cambridge, MA, MIT Press.

Moss HE and Gaskell MG (1999). Lexical semantic processing during speech. In Garrod, S & Pickering, M (Eds.) *Language processing.* Hove, UK, Psychology Press.

Moss HE and Marslen-Wilson WD (1993). Access to word meanings during spoken language comprehension: effects of sentential semantic context. *Journal of Experimental Psychology: Learning, Memory, and Cognition,* 19, 1254-1276.

Moss HE, McCormick SF and Tyler LK (1997a). The time-course of activation of semantic information during spoken word recognition. *Language and Cognitive Processes,* 12, 695-732.

Moss HE, Rodd JM, Stamatakis EA, Bright P and Tyler LK (2005). Anteromedial temporal cortex supports fine-grained differentiation among objects. *Cerebral Cortex,* 15, 616-627.

Moss HE and Tyler LK (2000). A progressive category-specific semantic deficit for non-living things. *Neuropsychologia,* 38, 60-82.

Moss HE and Tyler LK (2003). Weighing up the facts of category-specific semantic deficits: A reply to Caramazza & Mahon. *Trends in Cognitive Sciences,* 7, 480-481.

Moss HE, Tyler LK and Devlin J (2002). The emergence of category specific deficits in a distributed semantic system. In Forde, E & Humphreys, GW (Eds.) *Category-specificity in brain and mind.* Hove, UK, Psychology Press.

Moss HE, Tyler LK, Durrant-Peatfield M and Bunn EM (1998). 'Two eyes of a see-through': Impaired and intact semantic knowledge in a case of selective deficit for living things. *Neurocase,* 4**,** 291-310.

Moss HE, Tyler LK and Jennings F (1997b). When leopards lose their spots: Knowledge of visual properties in category-specific deficits for living things. *Cognitive Neuropsychology,* 14**,** 901-950.

Murphy GL and Medin DL (1985). The role of theories in conceptual coherence. *Psychological Review,* 92**,** 289-316.

Murphy GL and Wisniewski EJ (1989). Feature correlations in conceptual representations. In Tiberghien, G (Ed.) *Advances in cognitive science: theory and applications.* Chichester, Ellis Horwood.

Pecher D (2001). Perception is a two-way junction: Feedback semantics in word recognition. *Psychonomic Bulletin & Review,* 8**,** 545-551.

Pecher D, Zeelenberg R and Barsalou LW (2004). Sensorimotor simulations underlie conceptual representations: modality-specific effects of prior activation. *Psychonomic Bulletin & Review,* 11**,** 164-167.

Pexman PM, Holyk GG and Monfils MH (2003). Number-of-features effects and semantic processing. *Memory & Cognition,* 31**,** 842-855.

Pexman PM, Lupker SJ and Hino Y (2002). The impact of feedback semantics in visual word recognition: Number-of-features effects in lexical decision and naming tasks. *Psychonomic Bulletin & Review,* 9**,** 542-549.

Pilgrim LK, Moss HE and Tyler LK (2005). Semantic processing of living and non-living domains of knowledge across the cerebral hemisphere. *Brain and Language,* 94**,** 86-93.

Pitt D. (1999). In defence of definitions. *Philosophical Psychology,* 12, 139-159.

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

Price CJ and Humphreys GW (1989). The effects of surface detail on object categorisation and object naming. *Quarterly Journal of Experimental Psychology,* 41A**,** 797-828.

Pulvermüller F (1999). Words in the brain's language. *Behavioral and Brain Sciences,* 22**,** 253-279.

Randall B, Moss HE, Rodd JM, Greer M and Tyler LK (2004). Distinctiveness and correlation in conceptual structure: behavioral and computational studies. *Journal of Experimental Psychology: Learning, Memory, and Cognition,* 30**,** 393-406.

Raposo A, Stamatakis EA, Moss HE and Tyler LK (2004). Interactions between processing demands and conceptual structure in object recognition: an event-related fMRI study. *Journal of Cognitive Neuroscience,* 16**,** B82.

Rapp B and Caramazza A (1989). General to specific access to word meaning: A claim re-examined. *Cognitive Neuropsychology,* 6**,** 251-272.

Rips LJ (1989). Similarity, typicality, and categorization. In Vosniadou, S & Ortony, A (Eds.) *Similarity and analogical reasoning.* Cambridge, Cambridge University Press.

Rogers TT, Lambon Ralph MA, Garrard P*, et al.* (2004). Structure and deterioration of semantic memory: a neuropsychological and computational investigation. *Psychological Review,* 111**,** 205-235.

Rogers TT and Plaut DC (2002). Connectionist perspectives on category-specific deficits. In Forde, E & Humphreys, GW (Eds.) *Category-specificity in brain and mind.* New York, Psychology Press.

Rosch E and Mervis CB (1975). Family resemblances: studies in the internal structure of categories. *Cognitive Psychology,* 7**,** 573-605.

Rosch, E, Mervis, CB, Gray, WD, Johnson, DM and Boyes-Braem, P (1976). Basic objects in natural categories. *Cognitive Psychology,* 8**,** 382-439.

Sacchett C and Humphreys GW (1992). Calling a squirrel a squirrel but a canoe a wigwam: a category-specific deficit for artefactual objects and body parts. *Cognitive Neuropsychology,* 9**,** 73-86.

Saffran EM and Schwartz MF (1994). Of cabbages and things: Semantic memory from a neuropscyhological point of view - a tutorial review. In Umilta, C & Moscovitch, M (Eds.) *Attention and Performance.* Cambridge, MA, MIT Press.

Saffran EM and Sholl A. (1999). Clues to the functional and neural architecture of word meaning. In CM Brown and P. Hagoort (Eds.). *The Neurocognition of Language*. Oxford: Oxfored Universtiy Press, (pp241-271).

Saffran EM (2000). The organization of semantic memory: in support of a distributed model. *Brain and Language,* 71**,** 204-212.

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

Sartori G and Lombardi L (2004). Semantic relevance and semantic disorders. *Journal of Cognitive Neuroscience,* 16**,** 439-452.

Silveri MC and Di Betta AM (1997). Noun-verb dissociations in brain-damaged patients: Further evidence. *Neurocase,* 3**,** 477-488.

Simmons WK and Barsalou LW (2003). The similarity-in-topography principle: Reconciling theories of conceptual deficits. *Cognitive Neuropsychology,* 20**,** 451-486.

Smith EE and Medin DL (1981). *Categories and concepts,* Cambridge, MA, Harvard University Press.

Tabossi P (1988). Effects of context on the immediate interpretation of unambiguous nouns. *Journal of Experimental Psychology: Learning, Memory, and Cognition,* 14**,** 153-162.

Taylor KI, Moss HE and Tyler LK (in press). The conceptual structure account: a cognitive model of semantic memory and its neural instantiation. In Hart, J & Kraut, M (Eds.) *The neural basis of semantic memory.* Cambridge, Cambridge University Press.

Tippett LJ, McAuliffe S and Farah MJ (1995). Preservation of categorical knowledge in Alzheimer's disease: a computational account. *Memory,* 3**,** 519-534.

Tranel D, Logan CG, Frank RJ and Damasio AR (1997). Explaining category-related effects in the retrieval of conceptual and lexical knowledge for concrete entities: Operational and analysis of factors. *Neuropsychologia,* 35**,** 1329-1340.

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

Tyler LK, Bright P, Dick E*, et al.* (2003). Do semantic categories activate distinct cortical regions? Evidence for a distributed neural semantic system. *Cognitive Neuropsychology,* 20**,** 541-559.

Tyler LK and Moss HE (1997). Functional properties of concepts: Studies of normal and brain-damaged patients. *Cognitive Neuropsychology,* 14**,** 511-545.

Tyler LK and Moss HE (2001). Towards a distributed account of conceptual knowledge. *Trends in Cognitive Sciences,* 5**,** 244-252.

Tyler, LK, Moss, HE, Durrant-Peatfield, MR and Levy, JP (2000a). Conceptual structure and the structure of concepts: a distributed account of category-specific deficits. *Brain and Language,* 75**,** 195-231.

Tyler LK, Moss HE, Galpin A and Voice JK (2002). Activating meaning in time: The role of imageability and form-class. *Language and Cognitive Processes,* 17**,** 471-502.

Tyler LK, Stamatakis EA, Bright P*, et al.* (2004). Processing objects at different levels of specificity. *Journal of Cognitive Neuroscience,* 16**,** 351-362.

Tyler LK, Voice JK and Moss HE (2000b). The interaction of meaning and sound in spoken word recognition. *Psychonomic Bulletin & Review,* 7**,** 320-326.

Vigliocco G, Vinson DP, Lewis W and Garrett MF (2004). Representing the meanings of object and action words: the featural and unitary semantic space hypothesis. *Cognitive Psychology,* 48**,** 422-488.

Vinson DP and 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 DP, Vigliocco G, Cappa S and Siri S (2003). The breakdown of semantic knowledge: Insights from a statistical model of meaning representation. *Brain and Language,* 86**,** 347-365.

Warrington EK (1975). The selective impairments of semantic memory. *Quarterly Journal of Experimental Psychology,* 27**,** 635-657.

Warrington EK and McCarthy R (1983). Category-specific access dysphasia. *Brain,* 106**,** 859-878.

Warrington EK and McCarthy RA (1987). Categories of knowledge: further fractionations and an attempted integration. *Brain,* 110**,** 1273-1296.

Warrington EK and Shallice T (1984). Category specific semantic impairments. *Brain,* 107**,** 829-854.

Zingeser LB and Berndt RS (1990). Retrieval of nouns and verbs in agrammatism and anomia. *Brain and Language,* 39**,** 14-32.