BEHAVIORAL AND BRAIN SCIENCES (2001) **24,** 453–509
Printed in the United States of America

Hierarchies, similarity, and interactivity in object recognition: "Category-specific" neuropsychological deficits

Glyn W. Humphreys
Behavioural Brain Sciences Centre, School of Psychology, University of Birmingham, Birmingham B15 2TT, United Kingdom
g.w.humphreys@bham.ac.uk

Emer M. E. Forde
Institute of Psychology, Aston University, The Triangle, Birmingham, B4 7ET, United Kingdom
e.m.e.forde@aston.ac.uk

Abstract: Category-specific impairments of object recognition and naming are among the most intriguing disorders in neuropsychology, affecting the retrieval of knowledge about either living or nonliving things. They can give us insight into the nature of our representations of objects: Have we evolved different neural systems for recognizing different categories of object? What kinds of knowledge are important for recognizing particular objects? How does visual similarity within a category influence object recognition and representation? What is the nature of our semantic knowledge about different objects? We review the evidence on category-specific impairments, arguing that deficits even for one class of object (e.g., living things) cannot be accounted for in terms of a single information processing disorder across all patients; problems arise at contrasting loci in different patients. The same apparent pattern of impairment can be produced by damage to different loci. According to a new processing framework for object recognition and naming, the hierarchical interactive theory (HIT), we have a hierarchy of highly interactive stored representations. HIT explains the variety of patients in terms of (1) lesions at different levels of processing and (2) different forms of stored knowledge used both for particular tasks and for particular categories of object.

Keywords: category-specific deficits; functional imaging; hierarchical models; interactive activation models; neuropsychology; object recognition; perceptual and functional knowledge

1. Introduction

Perhaps some of the most interesting findings reported in the neuropsychological literature concern patients who can successfully recognize some categories of objects but not others. JBR, who suffered temporal lobe damage following herpes simplex encephalitis, was able to give fairly precise descriptions of nonliving artefacts, but could only give very impoverished responses for living things (Warrington & Shallice 1984). JBR described a compass as "tools for telling direction you are going" and a briefcase as "a small case used by students to carry papers," but when asked to describe a parrot he said "don't know" and he described a snail as "an insect animal." Category specific recognition impairments for living things have been documented on numerous occasions in the neuropsychology literature (Basso et al. 1988; De Renzi & Lucchelli 1994; Forde et al. 1997; Sartori & Job 1988; Sheridan & Humphreys 1993; Silveri & Gainotti 1988; Warrington & Shallice 1984), and there are also a few reports of patients with impairments for nonliving things (Cappa et al. 1998; Hillis & Caramazza 1991; Sacchett & Humphreys 1992; Warrington & McCarthy 1983; 1987; 1994). Recent summaries of relevant cases are provided by Caramazza (1998), Forde (1999), Forde and Humphreys (1999), and Saffran and Schwartz (1994). These cases raise general and important questions about the nature of our knowledge about objects, and its neural implementation. For example, does the brain represent knowledge of specific categories in discrete areas? Alternatively, does knowledge about categories take a distributed form, with particular forms of information "weighted" for the recognition of some but not other categories? Might even distributed knowledge be structured in some way, with (for example) item-specific perceptual knowledge being differentiated from forms of conceptual and contextual knowledge concerned with the relations between objects? Might inter-object contextual knowledge be separate from knowledge of the actions performed with objects? How might such structural differences in knowledge representation affect the recognition of particular classes of object? What might be the implications for understanding what is often termed "semantic memory"? This paper is concerned with these questions, addressed in the light of work on category-specific deficits for the living and the nonliving.[1]

© 2001 Cambridge University Press 0140-525X/01 $12.50
https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

1.1. Category-specific confoundings

Before we embark on discussing the category-specific deficits that emerge following brain damage, we need to ensure that any effects are indeed real and not an artefact arising from some other confounding variables. For instance, one possibility is that category-specific deficits simply reflect the background of the person pre-morbidly. People may have difficulty identifying animals because they knew little about them in the first place and consequently these items suffer most after brain damage; other people who knew little about tools may present with a category-specific impairment for nonliving things, and so on. A more elaborated account along these lines is that the deficits reflect differences in familiarity between objects. Living things are often less familiar (as a category) than nonliving things which we see and use every day, and so an apparent impairment in accessing knowledge about living things may be due to a more general impairment in retrieving information about low familiarity items (Funnell & Sheridan 1992). Another possibility is that object recognition is more difficult for more visually complex stimuli, taxing the limited visual processing resources of some patients (Stewart et al. 1992). Because living things are typically more visually complex than nonliving things, patients may present with what appears to be a category-specific impairment for living things.

Investigators should rightly be concerned that any effects they observe are not confounded by stimulus familiarity or complexity, and it is the case that many of the early studies did not control for such factors. Nevertheless, it seems unlikely that these factors can account for all the dissociations that have been observed. For example, some patients show category-specific losses for objects that they are particularly familiar with. We have observed a deficit in naming fruits and vegetables in a patient who was a food expert and wine connoisseur (Humphreys et al., in preparation), and Michelangelo, a patient with impoverished knowledge for living things, was formerly an active member of the World Wildlife Fund and could identify large numbers of animals before his brain damage (Sartori & Job 1988; Sartori et al. 1993a; 1993b). Also, the deficits apparent in patients can occur with objects that are highly familiar to average members of the population (e.g., apples, dogs, cats) so it is difficult to attribute impairments with these objects to a lack of familiarity. In addition, a simple effect of one variable (like familiarity or visual complexity) cannot account for the double dissociation between losses of knowledge for living and nonliving things, and such dissociations still occur even when items are matched for familiarity and complexity across categories or when statistical measures are taken to rule out such effects (Farah et al. 1996; Forde et al. 1997; Kurbat 1997; Kurbat & Farah 1998; Sartori et al. 1993a; 1993b). We conclude that not all such deficits are due to confounding factors, and thus the study of these category-specific impairments can help inform us about the nature of our stored knowledge and the way in which different objects are recognised.

Perhaps the most obvious account of category specific deficits is that they reflect the categorical organization of our underlying knowledge about the world, which is partitioned according to whether stimuli are living or nonliving. We will use this as the "default" account of the deficit, to be held at the back of the reader's mind while other "non-categorical" accounts are reviewed, though we will reconsider it following these reviews. Prior to this, we will evaluate whether the deficits reflect the loss of particular forms of knowledge that are not categorical in nature (sect. 2) or whether they reflect the interaction between perceptual processes and the kinds of knowledge required to differentiate between objects for different tasks (sect. 3). In section 4, we return to the idea that stored knowledge might be categorically organized. In section 5, we consider all the above arguments in relation to function imaging studies that have shown selective activation of neural areas functional according to both the object being presented and the task. In section 6, we outline a framework for understanding both the psychological and the neuro-anatomical data which we term the hierarchical interactive theory (HIT) of object recognition and naming. The model is related to several accounts that have previously been considered (e.g., Damasio 1989; 1990; Warrington & McCarthy 1987), but differs in emphasising that object recognition proceeds through a hierarchical series of stages and that re-entrant activation plays a differential role in particular tasks (e.g., for naming rather than for object recognition). We discuss the relations between the model and others in the literature.

Glyn W. Humphreys is Professor of Cognitive Psychology at the University of Birmingham. He obtained his BSc. and PhD. from the University of Bristol and was Lecturer, Senior Lecturer, and then Professor of Psychology at Birkbeck College before moving to Birmingham in 1989. He has worked extensively on experimental and neuropsychological studies of visual object recognition, word recognition, attention and action, with this work extending to computational modelling and functional anatomical studies of visual cognition. He has been awarded both the Spearman Medal and the President's Award for research by the British Psychological Society. He is currently the editor of *Visual Cognition*.

Emer M. E. Forde is a Lecturer in Psychology at the Neurosciences Research Institute, Aston University. After graduating from Cambridge University in 1992, she obtained her PhD. and worked as a postdoctoral research fellow at the University of Birmingham before moving to Aston. In addition to neuropsychological work on category-specific impairments, semantic memory, and everyday actions, she is currently investigating the brain areas involved in the recognition of living and nonliving things using magnetoencephalography. Emer Forde and Glyn Humphreys are authors of *Category-specificity in brain and mind* (Psychology Press, 2000).

1.2. Defining semantic memory

Category-specific impairments have frequently been interpreted in terms of deficits in accessing semantic memory for objects. Consequently, such deficits have been used to argue about the nature of semantic memory (e.g., Warrington & Shallice 1984). But what is semantic memory? In Warrington's (1975) seminal paper, which was the first systematic investigation of patients with selective impairments of long-term stored knowledge about objects and their properties, the term "semantic memory" was defined as "that system which processes, stores and retrieves information about the meaning of words, concepts and facts." Consistent with this, semantic memory is still considered by

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

many to hold the information that allows us to give meaning to the objects we see and the words we read and hear. However, despite the general use of the term, or perhaps because of it, there has been little attempt to provide a more rigorous definition (though, see Caramazza et al. 1990). It is perhaps symbolic that semantic memory is represented as an under-specified "cloud" in many standard models of cognition (e.g., see Morton & Patterson 1980, Figs. 4.1–4.3). Now, our semantic representation of a dog would generally be considered to include multiple facts such as: it is an animal, it has four legs and a tail, it barks, it likes to chase cats, it is a "man's best friend," it was once featured in a sentimental song by Elvis Presley and so on. Thus, on this basis, semantic knowledge includes information about the general category of the item, visual information about its shape and parts, other sensory information (e.g., about the sound it makes, what it feels like), the relationship between it and other items, and general contextual knowledge abstracted even from any sensory information of the object itself (e.g., that the ballad "Old Shep" refers to a dog). Given the diversity of the information considered to be "semantic," and the fact that there are fundamental differences in the nature of the knowledge involved (i.e., some pertains to the sensory properties associated with an object [e.g., it is brown and has floppy ears] and some reflects how two or more objects relate to each other [e.g., dogs chase cats], etc.), it is perhaps surprising that all of this information is considered to be represented in one homogenous "store." Indeed, even when attempts have been made to define subsets of semantic knowledge, this has typically led to dichotomous distinctions being made, such as between visual and verbal semantics (see Riddoch et al. 1988, for a review). One aim of this paper, through the review of category-specific deficits, is to argue that the concept of a "semantic system," in any unitary sense, may be one of the victims of an attempt to define the nature of our stored knowledge in more detail. If the "semantic system" retracts to no more than the form of knowledge recruited to perform a particular task – with this knowledge differing across tasks, then the "system" becomes a fiction. For instance, there may be little more in common between different forms of semantic knowledge than there is between, say, different forms of sensory knowledge (visual, auditory). It may be more fruitful to specify the different forms of knowledge than to seek out some unifying principle across what turn out to be separable knowledge stores. However, because the notion of a unified semantic system has been at the heart of much of the research, we set out by using this term when we discuss the work. In this initial usage, semantic memory may be defined as the central knowledge store for all input and output modalities that contains information about the meaning of objects – much like a multi-modal dictionary that can both be accessed and can express itself in a variety of different ways (from print, speech, and visual images as input, to print, speech, and action as output).

2. The loss of particular forms of knowledge: Sensory and functional knowledge

2.1. The sensory/functional distinction

Warrington and Shallice (1984) reported four case studies of patients who had particular problems in identifying living things. For example, the patient we described in the Introduction, JBR, was able to give precise definitions of nonliving things (compass and briefcase) but very poor definitions of living things (parrot and snail). In addition to his poor definitions of living things, JBR was only able to produce descriptions indicating identification for 6% of pictures of living things, but 90% of nonliving things. A second patient, SBY, named no pictures of living things but could produce correct descriptions for 75% of the pictures of nonliving things. Warrington and Shallice suggested that since the patients were impaired at accessing information from both words and pictures, the locus of their impairment must be within the semantic system, rather than a lower level visual recognition problem. They also noted that their patients were poor at defining some nonliving things, such as cloths and precious stones. An associated problem with these particular nonliving things is difficult to understand if the patients' problems are confined to living things. Instead, Warrington and Shallice proposed that the deficits reflect loss of some but not all forms of semantic knowledge about objects – notably loss of sensory semantic knowledge. They argued that, in order to identify living things (fruit, vegetables or animals), retrieval of fine-grained sensory information was necessary. For example, they suggested that to distinguish between a raspberry and a strawberry, detailed information about color, size, shape, and texture was necessary. In contrast, they suggested that recognition of nonliving things "depends crucially on determination of its functional significance" (p. 849). What is actually meant here by functional significance is not clear. Our knowledge about the function of an object could include information about how an object is acted upon (e.g., by turning the wrist, for a spanner) or about how the object itself operates (e.g., a car operates by consuming petrol). As we hope to show, authors have used the term "functional" in a variety of ways when referring to our stored knowledge of objects. The different forms of functional information, however, need not be equivalent, and may themselves be represented in contrasting ways.

Warrington and Shallice proposed that two independent systems may have evolved: one storing "functional" information important for identifying nonliving objects, and the other storing sensory information important for identifying living things. They suggested that patients like JBR and SBY have an impairment to the visual/perceptual semantic system which leads to particular problems naming living things (see Fig. 1). However, since the source of the problem is in a store specifying sensory rather than category-specific knowledge, the problem can also generalize to nonliving things that also depend on the retrieval of sensory knowledge for their identification – cloths and precious stones perhaps being two examples.

Warrington and Shallice also noted that their patients tended to be either consistently correct or consistently incorrect when trying to identify stimuli and suggested that such a consistent deficit reflected degenerate semantic knowledge. Interesting to note, the items that the patients were consistent on in one modality (say, with pictures) were not necessarily the same as those for which they were consistent in other modalities (say, when defining a word). From this they concluded that semantic knowledge is partitioned not only into independent modules for sensory and functional knowledge, but also for input modality. Thus there may exist sensory semantics for visual objects (pic-

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

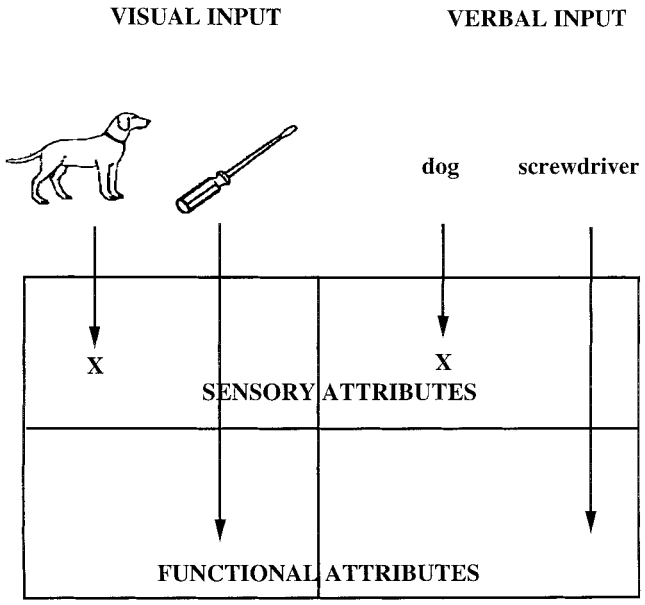


Figure 1. A model which distinguishes knowledge stores for sensory and functional knowledge, separately for visual and verbal input (after Warrington & Shallice 1984). Note that, according to this account, patient JBR (Warrington & Shallice 1984) – who has problems with living things when given both visual and verbal input – has functionally independent impairments in the visual/sensory and verbal/sensory semantic systems.

tures and real objects) and sensory semantics for verbal input (written and spoken words). Since these semantic systems are distinguished by input modality as well as by the information represented, there may well be duplication of sensory and functional knowledge within the modality-specific systems (e.g., the fact that a bee has yellow and black stripes will be represented in both the visual and verbal semantic systems specifying the sensory properties of objects) (see Fig. 1; also McCarthy & Warrington 1994; Warrington & McCarthy 1994). We return to this issue of modality in the general discussion (sect. 6.5).

Warrington and McCarthy (1987) presented a rather different view of this sensory/functional dichotomy, which they discussed in terms of a distributed model of semantic memory. They drew upon ideas of how objects might be recognized put forward originally by Lissauer (1890). Lissauer suggested that

> the recognition of an object can only occur when at the time of its perception a number of ideas are evoked which relate to that object. These bring into consciousness those characteristics which the mind has learned to associate with it and those conditions in which it has been previously experienced . . . memories laid down through different sensory modalities contribute to these associations but it is only when they are brought into awareness and linked with the percept that the recognition of an object becomes complete.

Warrington and McCarthy (1987) suggested that these "memories laid down through different sensory modalities" would not be stored in a homogenous semantic store or "module" but in modality congruent "channels." For example, visual information would be stored in a visual channel, "functional" information in a motor channel, information about the sound an object makes in an auditory channel, and so on. Warrington and McCarthy suggested that these channels could be relatively fine-grained so that visual information might actually be stored in a number of subchannels (e.g., for colour, size, shape). Furthermore, different channels of sensory or motor information would have different degrees of importance for different items. For example, as initially proposed by Warrington and Shallice (1984), perceptual/sensory information would be important for the identification of many living things whereas "functional" information, defined to include motor actions, may be crucial for the identification of nonliving things. The specialization of different parts of the system for particular objects would lead to a quasi-categorically organised knowledge base.

This account predicts that quite fine-grained category-specific impairments may occur because different categories within the living and nonliving groups (e.g., animals, tools, fruit, clothing) would have different patterns of weighting across the channels. For example, they suggested that accessing colour knowledge might be important for recognizing fruit (e.g., distinguishing between a raspberry and a blackberry) but accessing shape information may be relatively more important for distinguishing between two flowers (e.g., a daffodil and a tulip). Thus patients may have deficits with subsets of living or nonliving things, for example, with fruit and vegetables but not animals, or with tools but not clothing. These more selective patterns of deficit have been observed (e.g., Hart et al. 1985, for a deficit with fruit and vegetables; Caramazza & Shelton 1998 and Hart & Gordon 1992, for a deficit with animals; McCarthy & Warrington 1987, for a deficit with tools). Also, patterns of association are expected between the loss of certain forms of knowledge and deficits with particular objects. For example, if a patient had a specific impairment in naming fruit, this ought to be accompanied by an impairment in retrieving the colour of objects from memory, since colour knowledge is likely to be important for distinguishing between different fruit (see Price & Humphreys 1989, for evidence with normal subjects). However, Luzzatti and Davidoff (1994) reported two case studies of patients who had a marked impairment at retrieving the color of objects, but no particular problem in naming fruit and vegetables. Luzzatti and Davidoff argued that an impairment in retrieving object-colour knowledge did not necessarily impair naming performance for categories of living things, such as fruit and vegetables.

An attempt to capture some of the properties of the sensory-functional distinction in a more formal model was made by Farah and McClelland (1991). They simulated semantic representations using a distributed associative memory system (cf. McClelland & Rumelhart 1985). Items were represented in terms of patterns of activation across processing units corresponding to either the perceptual-sensory or functional properties of objects. The number of sensory units assigned a non-zero value in the coding of a stimulus, relative to the number of functional units, differed for living and nonliving things. For living things, about seven times more sensory than functional features were active; for nonliving things, there were equal proportions of active sensory and functional units. This differential weighting of the representations of living things was based on the number of properties generated by subjects when asked to mark the visual and functional attributes in dictionary definitions, and far more visual attributes than functional attributes were marked. In contrast, roughly

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

equal numbers of visual and functional attributes were marked for nonliving things. The model was then trained to associate the sensory and functional properties of objects with representations of their visual attributes and names. Performance on object naming was tested by giving the visual attributes as input and seeing whether the correct name was generated as output. The model was also lesioned, with sensory or functional units being differentially affected. Lesions affecting the sensory units lead to marked impairments in identifying living things; lesions affecting the functional units affected non-living things more than living things. In addition, because of the distributed nature of the semantic system, lesions of the sensory units also produced some loss in retrieving the functional properties of stimuli, with living things being affected more severely.

Farah and McClelland's simulation provides an existence proof that a model of this form, with distributed sensory and functional knowledge about objects, can produce apparent category-specific deficits when lesioned. It also shows that a form of double dissociation can be generated, with either living or nonliving things being affected – depending on which form of knowledge is impaired. We consider this argument further in our discussion of category-specific deficits for nonliving things. However, such an existence proof does not demonstrate that the same underlying architecture exists for the stored knowledge of objects in humans. Also, Farah and McClelland's procedure for estimating the sensory and functional properties of objects has been criticised (see Caramazza & Shelton 1998). Farah and McClelland asked subjects to mark the functional attributes of objects by asking the questions: "What does the item do or what is it used for?" Caramazza and Shelton suggest that this question is biased to nonliving things and leads to underestimates of what might be termed the functional properties of living things (here Caramazza & Shelton refer to nonsensory properties such as "lives in a desert," and "carnivore"). When all nonsensory properties are noted in dictionary definitions then the bias for more sensory properties for living things is greatly reduced (Caramazza & Shelton 1998; Moss et al., in press).

In the distributed associative memory framework used by Farah and McClelland, category-specific deficits for living things should generalise to include poor retrieval of the nonsensory properties of these items (although this problem should be less severe than the impairment in retrieving sensory properties). We run into two difficulties. First, there are patients whose problems in the retrieval of nonsensory attributes of living things is at least as severe as their problems in retrieving sensory properties (Caramazza & Shelton 1998; Laiacona et al. 1993; Sheridan & Humphreys 1993). Indeed, this pattern of an equal deficit for sensory and nonsensory knowledge of living things was observed in one of the original Warrington and Shallice cases, when reexamined by Funnell and de Mornay-Davies (1996). Second, there are some patients who show extremely good (normal) retrieval of the nonsensory properties of living things whilst being impaired at retrieving their sensory properties (Forde et al. 1997; Hart & Gordon 1992; Humphreys et al. 1997; Riddoch & Humphreys 1993). This pattern, in which nonsensory knowledge can be affected to as severe a degree, to a less severe degree or not at all, suggests that the sensory and non-sensory ("functional") properties of living things are not as tightly coupled as in Farah and McClelland's architecture. The results also have implications for accounts of category-specific deficits in terms of correlated sensory and functional features (sect. 2.2) and in terms of a semantic system that is truly categorical in nature (Caramazza & Shelton 1998) (sect. 4).

2.2. Correlated sensory and functional features

We have suggested that a simple distributed memory account, which separates sensory from non-sensory knowledge fails to capture the full pattern of deficits for these two forms in patients. Other investigators have suggested that the difference between living and nonliving categories lies not in the relative importance of perceptual or functional attributes for identification per se, but in the links between them. For example, for nonliving things, the connection between shape and action is not arbitrary, since the shape of the item is typically constructed in a way that will best perform the action intended. De Renzi and Lucchelli (1994) argued that this link between visual and functional (in the sense of action) properties for nonliving things makes these items less vulnerable when stored sensory information is degraded. They outlined a case study of a patient who had difficulty in recognising living things and in performing object decisions, drawing and describing the perceptual differences between living things. Also interesting, their patient also had problems with nonliving things when the tasks involved retrieving the colour of items and naming objects from sound. De Renzi and Lucchelli proposed that the naming impairment for living things resulted from a general failure to retrieve the perceptual features of objects from every category (and not just the category of living things). However, for nonliving things this deficit may be compensated for by the close links between visual attributes and function, which provide an alternative route to accessing the representation of the object. De Renzi and Lucchelli acknowledged that "it remains to be explained *how* the function of an object can be inferred from its visual appearance, if this has not been recognized or retrieved from memory" (p. 19, emphasis ours). However, they suggested that for nonliving things, a sensory and a functional semantic store could interact. Thus, in visual object identification, functional cues could "help specify hypotheses on the nature of the stimulus that were left undefined by visual processing" (p. 20). For living things they proposed that this route would be unavailable because, for these items, they assumed there to be few links between perceptual and functional attributes. In addition, they suggested that the retrieval of colour knowledge and identification via sound would be impaired for all categories because these perceptual properties have no direct functional associations.

Another account that stresses the importance of the interaction between sensory and functional properties of objects is the OUCH model of semantic memory (Caramazza et al. 1990). This model does not differentiate between different types of stored knowledge; both perceptual and functional (action-related) properties are said to be stored within a single semantic system. OUCH states that, during visual object processing, salient parts of objects directly activate corresponding perceptual and functional attributes in semantic memory. For example, salient parts of a fork (e.g., the tines and handle) directly activate corresponding semantic knowledge (i.e., about tines and handles in general) and this pattern of activation could then be used to gener-

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

ate hypothesis about what the object might be. This direct activation of action-related information from the parts of the object leads to privileged access to semantic memory for objects relative to words (cf. Potter & Faulconer 1975). According to this account, damage to semantic memory may lead to a more severe impairment for living relative to nonliving things, if nonliving things can benefit from a higher degree of correlation between their sensory and functional features.

Both De Renzi and Lucchelli's account and the OUCH model emphasise the importance of correlations between sensory and functional properties of objects. Other authors, however, stress that the correlations within sets of sensory and/or sets of functional features may also play a predictive role in retrieving information from semantic memory. McRae et al. (1997) had subjects list the features of objects from categories of both living and nonliving things. They found that, whilst sensory properties were listed with roughly equal frequency for living and nonliving things, there were significantly more nonsensory features (e.g., used for carpentry, worn by women) for nonliving things (supporting Warrington & Shallice's [1984] model of semantic memory). Furthermore, living things tended to have more correlated features than nonliving things (11% of feature pairs, relative to only 6% of feature pairs for nonliving things). Living things were more densely represented across the correlated feature pairs – so that a smaller set of common features captured more of the properties of living things. They suggested that correlated features play a differential role in recovering information about living and nonliving things; for living things, correlated features could lead to robust recovery of a common set of core attributes (e.g., is animate, eats, breathes) but greater difficulties in individuating objects. This proposal can account for some patterns of degenerative performance found in patients with Alzheimer's disease (AD). Gonnerman et al. (1997) reported that in the early stages of AD, living things could be identified better than nonliving things, but as the disease progresses the pattern reverses and living things are relatively more difficult to identify (see also Devlin et al. 1998). The greater number of correlated features for living things may help to protect the identification of these items from small amounts of generalised brain atrophy (as in the early stages of AD). With more widespread damage though, features are lost and performance with living things may decrease catastrophically, as the lost features contribute to many exemplars within the category. Nonliving things, having relatively fewer correlated features, are vulnerable to damage on a more individual basis, but then show a less catastrophic loss as the disease progresses. This differential pattern of decline was simulated by Devlin et al. using a neural network with interconnected semantic units that tended to "push" an input pattern into a stable activation state. Interestingly, in this model, category specific deficits emerged even with random non-selective damage – as might be assumed to occur in AD (though see Perry 1999, for a discussion of the limitations of the model). In models such as that of Farah and McClelland (1991) lesions had to affect either the sensory or functional features differentially to generate a category-specific effect.

However, the data on patients with AD are not clear cut. For example, Silveri et al. (1991) found that AD patients with moderate deficits were worse with living than with nonliving things (see also Garrard et al. 1998; Guistolisi et al. 1993) – the opposite result to that reported by Devlin et al. (1998) and Gonnerman et al. (1997). Gonnerman et al. in fact found a consistent deficit for living things in one patient even when identification performance for nonliving things remained at a high level, which is not consistent with their group study. It is not clear that accounts in terms of correlated features will be able to provide a framework for these disparate results. Garrard et al. (1998) also query how, in neural terms, inter-correlations between representations might protect features from degenerative decay. They suggest that the differences between the majority of degenerative patients who have problems with living things, and the minority who can be found with deficits for nonliving things, related to differences in the initial area of neocortical involvement. In the majority of cases there is transfer of the disease from the transentorhinal region to temporal neocortex. In the minority there may be bi-parietal involvement. The contrast between the two sets of patients may reflect the storage of perceptual features in temporal cortex and more action-based (functional) features in frontoparietal regions (see also Gainotti et al. 1995).

An extension to the proposal that correlated features are important has been made by Moss, Tyler, and colleagues (Durrant-Peatfield et al. 1997; Moss et al. 1997; 1998; Tyler & Moss 1997). They point out that, for living things, many of the correlated perceptual features are associated with common biological functions (such as breathing, eating, and reproducing). The distinctive perceptual features of living things, however, are not strongly correlated with this kind of functional information (e.g. a tiger's stripes; see Keil 1987). It is in this last respect that living and nonliving things differ. Nonliving things have distinctive perceptual features that are correlated with their function in terms of action (e.g., the serrated edge of a saw, the tines of a fork). This last point is similar to the argument made by De Renzi and Lucchelli and the OUCH model (see above), the difference being that Moss, Tyler, and colleagues highlight that, for living things, common (intercorrelated) perceptual features are associated with functional properties whilst for nonliving things, functional attributes are associated with distinctive perceptual properties. After brain damage, the features strongly associated to the functional properties of objects may be better preserved than those with weak associations. This will have different consequences for living and nonliving things. For living things, information about biological function will be recovered from the linked, common perceptual features. However, this will not help the identification of individual stimuli, for which distinctive features are important. For nonliving things, distinctive visual features are correlated with function; identification of individual items is thus better. Durrant-Peatfield et al. (1997) report simulations of these patterns in a feedforward connectionist model trained to associate an input pattern corresponding to "perceptual" and "functional" properties of objects to a matching output pattern ("auto-associative" learning). The training set was varied so that, for living things, common perceptual features co-occurred with common functional properties; for nonliving things, distinctive perceptual features co-occurred with distinctive functional properties. When connections between input and output units were randomly disconnected, the identification of living things tended to be more affected, though recovery of their shared functional properties was relatively well preserved.

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

The account put forward by Moss et al. is able to explain why information about biological function can be preserved in patients with poor identification of living things (Moss et al. 1997; Tyler & Moss 1997). For example, in one case reported by Moss et al. (1998) a patient with impaired naming of living things was nevertheless able to group these items according to their shared properties (does it have legs? Does it lay eggs?). He was poor at retrieving distinctive properties of living things, whether visual or functional (e.g., properties concerned with survival). Retrieval of shared category properties for nonliving things was, if anything, worse than retrieval of similar properties for living things, perhaps because of an inherent advantage for living things in such tasks (due to the shared functional properties being based on shared perceptual properties, for living things; see sect. 3.1). Indeed common information for living things, whether concerned with biological function or general category, is accessed rapidly also by normal subjects (Humphreys et al. 1997; Tyler & Moss 1997).

We conclude that differences between shared and distinctive features, and the degree to which these features correlate with the function of the object, are likely to be important contributing factors in category specific impairments. However, as we hope to demonstrate, models need to be elaborated further in order to account for the full pattern of dissociations that have been documented. In particular, models need to have a more articulated structure, specifying different forms of stored knowledge and different stages of object identification. HIT, which differentiates both the different forms of knowledge representation and the contrasting stages of object identification, provides a framework that can allow for a fuller account of the different patients in the literature. This is described in section 6.

3. Interactions between perceptual processes and knowledge for particular tasks

3.1. The Cascade model

When normal subjects are asked to list the parts of objects, living things tend to be listed as having proportionately more shared parts than nonliving things. Similarly, when the outline contours of standardised drawings of objects are compared across category exemplars, living things tend to have higher levels of contour overlap than do nonliving things (see Humphreys et al. 1988). These different indices provide an approximate measure of the similarities of the perceptual structure of objects within their categories, with living things having more similar structures than nonliving things. Humphreys et al. (1988) termed this within-category property "structural similarity." Differences in structural similarity between living and nonliving things may contribute to the differences in identification that can be observed between these categories of object. To illustrate, consider the "Cascade" model of visual object recognition outlined by Humphreys et al. (1988). This is composed of several stages including visual recognition of an object's structure (access to stored structural descriptions), access to semantic information, and access to the object's name. Stored structural descriptions are held to represent information about the shape of objects but not to include other information such as an object's use or its association with other objects.[2] In this case, the term "semantic memory" was reserved to apply only to the latter forms of (non-perceptual) knowledge. If activation can be passed on to one stage before processing at an earlier stage is completed (i.e., if processing operates in cascade), then differences in structural similarity will directly affect semantic access and name retrieval when objects are presented visually. Structurally similar objects will activate the structural representations of perceptual neighbours across their category. As a consequence, functional and associative information common to the category is derived quickly, but there is then increased competition between category exemplars for individual identification. Structurally dissimilar objects will activate fewer perceptual neighbours, so that activation of functional and associative information will be slower and less widespread. Nevertheless, individual identification should be more efficient (e.g., in a naming task), since competition from perceptually and functionally similar neighbours will be reduced (see Humphreys et al. 1988; 1997). Note that this account of rapid access to common functional and associative information, along with slowed access to identity information owing to within-category structural similarity for living things, is in many ways similar to the proposals concerning common and distinctive features made by Moss, Tyler, and colleagues. In the Cascade model, though, these ideas are tied to an architecture specifying the different stages involved in object identification. We will argue that this provides important explanatory power in accounting for the variety of category specific deficits in patients, and it remains a feature of the HIT approach which we outline in the final section.

The Cascade model predicts differences in performance for living and nonliving things even in normal subjects. Consistent with this, Humphreys et al. (1988) found that normal subjects named pictures of living things more slowly than pictures of nonliving things matched for familiarity and name frequency (see also Lloyd-Jones & Humphreys 1997; Snodgrass & Yuditsky 1996). When access to stored structural descriptions is measured using object decision, the benefit for nonliving things remains but is reduced (Lloyd-Jones & Humphreys 1997). According to the model, the larger difference between living and nonliving things in naming, relative to object decision, is a result of the small differences in the efficiency of accessing structural descriptions being exacerbated by competition accruing from common (overlapping) functional and associative representations being activated. However, category decisions are faster for living things compared to nonliving things (Humphreys et al. 1999; Riddoch & Humphreys 1987a). This advantage for living things is larger when stimuli are presented as pictures rather than words, though it still exists with words (Job et al. 1992). Since the effect is larger with pictures, we suggest that it is not simply due to category information being more closely linked to living things; rather there is privileged access to common semantic (functional and associative) information from the visual properties of living relative to nonliving objects. Advantages for identifying nonliving things over living things in fact occur not only with normal human subjects but also with monkeys! Gaffan and Heywood (1993) trained monkeys to make discrimination responses to pictures of living and nonliving things. They found that the monkeys took longer to learn the responses associated with living things, with the effect increasing as the number of stimuli in the set increased. This work provides converging evidence for the idea that living things have high levels of perceptual overlap.

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

A simulation of the results on human identification was reported by Humphreys et al. (1995). They used an interactive activation and competition framework with pools of units representing structural, semantic (functional and associative) and name information about objects (see Fig. 2). Name units were further divided to represent either specific or superordinate category names. Input activations, given to the structural descriptions, were based on the rated perceptual similarity between items within their categories; for example, a "dog" as input would maximally activate its own structural description (activation 1), but it would also activate the structural descriptions of other, similar items (e.g., activating descriptions for a fox, a cat, a sheep, with values of .6, .5, .3, etc.). Living things were rated as perceptually more similar than nonliving things, and so shared more activation values between their category members. The net effect of this difference in shared activation values was that there was rapid access to superordinate names for living things but slowed access to individual names (due to increased competition from multiple activations at a semantic level). This model was also "lesioned" by having noise added either to activations at a structural description level or to the weights connecting structural to semantic, or semantic to name representations. After lesioning, the identification of living things was selectively disrupted relative to the identification of nonliving things; essentially this occurred because of the increased structural and functional/associative overlap generated during the visual processing of living things. In contrast to the effects on identification there was relatively little effect of lesioning on access to superordinate information; indeed, if anything, the disruption affected performance with nonliving things more. These data from the lesioned model match the results from many of the patients with category-specific deficits with living things.

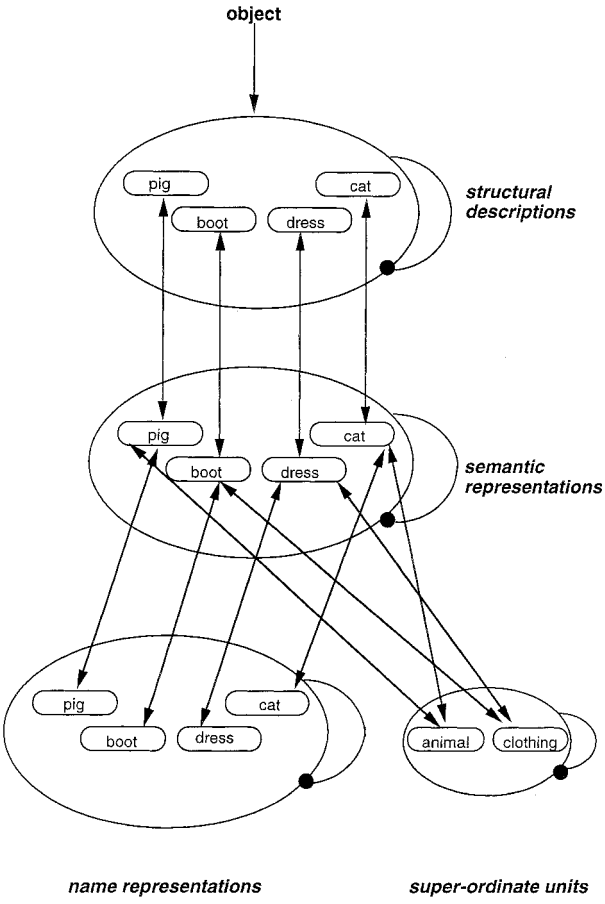


Figure 2. The interactive activation and competition model of object naming proposed by Humphreys et al. (1995). ⟶ indicates excitatory links and —● inhibitory links. This model incorporates two of the basic principles of the HIT framework: functionally isolatable subsystems (e.g., separating perceptual knowledge about objects [in the structural description system] from other forms of stored knowledge), and interactive processing (top-down as well as bottom-up). In a more detailed account, the semantic representations specified here are further divided and represented in a distributed form.

One of the interesting results in these studies of simulated lesions was that the identification of living things was impaired even when noise was only added to the connections between the semantic and name units (i.e., at the "output" end of the model). In this last case, the impairment still arose even when access to structural and semantic information was intact (though there was still the natural increase in competition for access to the names of living things). This result is of interest because it matches the pattern reported in a few patients with category specific impairments for living things. For example, Farah and Wallace (1992) and Hart et al. (1985) have both reported patients with problems that seem confined to name retrieval. In both cases, the patients could retrieve information about living things once they were given their names, and Farah and Wallace report that their patient could categorise fruit and vegetables even though the naming of these items was selectively impaired (see also Forde et al. 1997; Humphreys et al. 1997). Such an apparent selective problem in name retrieval is difficult to account for in models that do not have a distinct (and dissociable) stage of access to name information following access to semantics. Nevertheless, despite the success of the Cascade model in accounting for such naming disorders, we shall reconsider the evidence on this point in section 6, when we introduce the HIT model of object identification.

3.2. Depth of processing within the structural description system

The Cascade model predicts that category specific impairments for living things can emerge following damage to the structural description system, though this description system itself is not categorically organised. According to the model, damage to structural representations may lead to problems in differentiating between living things because these stimuli belong to categories with large numbers of structurally similar exemplars. Category-specific impairments at the level of the structural description do not reflect any functional distinction based on biological category but emerge because of the different processing demands of living and nonliving things. However, Sartori and Job (1988) offered an alternative suggestion and proposed that the structural description system might be categorically organised. They presented a case study of a patient, Michelangelo, who was significantly worse at naming living compared to nonliving things. In addition, Michelangelo performed poorly on object decision for living things, though performance with nonliving things was within normal limits. Object decision is typically used to test the integrity of the

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

structural description system, because it only requires subjects to assess whether or not they have seen a particular "visual pattern" before and does not require them to access any further (functional or phonological information) information. Consequently, Sartori and Job (1988) suggested that Michelangelo's category-specific impairment resulted from damage to the structural description system. They suggested further that the particular problem with living things could arise because (1) the system was categorically organised or (2) living things require more detailed processing at this stage within the object recognition system (as suggested in the Cascade model, Humphreys et al. 1988). To distinguish between these two alternatives they asked their patient to specify the important differences between 12 pairs of animals, 12 pairs of vegetables, and 12 pairs of objects. Michelangelo scored at ceiling for objects but outside the control range for vegetables and animals.

Sartori and Job suggested that if Michelangelo had a general impairment in accessing perceptual attributes from stored structural descriptions he should have problems generating the important perceptual differences between objects, regardless of category. Since he only had problems with living things, they argued that these must be represented in a functionally independent compartment. However, if living and nonliving things are represented in terms of contrasting amounts of overlapping activation within the structural description system, then we see no reason why differences should not emerge between the categories even when the system is addressed from the spoken names. Increased overlap for living things should still make their perceptual properties relatively difficult to retrieve. Furthermore, both Michelangelo and controls sometimes gave perceptual attributes and sometimes gave functional attributes for items from all categories, when contrasting the differences between stimuli. Controls tended to give functional attributes for objects and perceptual attributes for animals and vegetables. Consequently, if Michelangelo had intact functional information but poor perceptual information for all categories, he would be able to score more highly with objects compared to the other two categories, as Sartori and Job observed. Stronger evidence for the idea that the structural description system is categorically organised would be reports of patients who have impairments to the structural description system that selectively affect nonliving things; this pattern remains to be documented.

More recently, Sartori et al. (1993) have used Marr's (1982) model of visual object recognition to account for category-specific deficits for living things. According to this model, structural descriptions of objects are represented in a hierarchical form. At the top of the hierarchy, descriptions correspond to general category exemplars, with the descriptions becoming increasingly detailed further down the hierarchy, eventually distinguishing items at a subordinate level. Sartori et al. suggested that patients with category-specific impairments for living things may have problems accessing the lowest (most detailed) levels of these structural hierarchies, and argued that living things were most affected because they have "deeper" representations than nonliving things. Note that this idea is similar to the view outlined by Humphreys et al. (1988), who also argued that living things require more fine-grained or "deeper" processing at the level of structural descriptions. For Humphreys et al. (1988), "deeper" processing is an emergent consequence of overlap between the structural descriptions of stimuli.

3.3. Shape processing differences across categories

The Cascade model emphasizes the importance of perceptual overlap between exemplars within a category as a factor that can lead to processing differences between living and nonliving things. However, the approach suffers from a failure to specify in detail the nature of the perceptual overlap that may be critical. Attempts to overcome this have been made by Arguin, Bub, and colleagues (e.g., Arguin et al. 1996; Dixon et al. 1997). They used computer generated stimuli, derived from variations in elongation, tapering, and bending, to resemble real objects. For example, a banana can be described as having positive values on elongation and bending, but a zero value on tapering; an orange would have zero values on all of these dimensions. Arguin et al. reported data from a patient, ELM, who showed a category-specific deficit for living things. One task involved presentation of four shapes, one in each quadrant of a display, followed by presentation of one of the shapes at the centre. ELM had to point to which of the four locations had been occupied by the subsequent target. When the items on a trial differed along a single dimension the patient performed significantly better (29% errors) than when the items varied along two dimensions simultaneously in the conjunction condition (57% errors). Arguin et al. proposed that ELM failed to extract information from two visual dimensions simultaneously, and tended instead to attend to just one dimension (so performing better when stimuli only varied along a single dimension).

Arguin et al. went on to show that this deficit in extracting multiple visual dimensions combined with effects of semantic similarity between items. ELM was required to label the same shapes either as particular fruit and vegetables or (in another block) as nonliving things. With the semantically close living things, ELM again performed worse in a conjunction condition relative to a single dimension condition. However, with a set of unrelated nonliving things, there was no difference between performance with conjunction and single feature sets. The deficit for learning conjunction stimuli paired with semantically close living items also generalized to faces (see Dixon et al. 1998). Arguin et al. argued that visual processing deficits, concerned with whether patients can extract several visual dimensions simultaneously, can combine with semantic similarity between items to create problems in identifying particular objects. To link these problems to a deficit with previously known living things, we must assume that living things, more than nonliving things, share both semantic and visual features; in addition, the visual features of living things may also vary along more than one dimension simultaneously. Support for this first assumption comes from data on normal picture naming under deadline conditions. Normal subjects make a greater range of errors that are both semantically and visually related to targets for living things than for nonliving things (Vitkovitch et al. 1993). Greater overlap in access to semantic as well as visual representations is also supposed by accounts that point out the importance of correlated perceptual features for living things (Gonnerman et al. 1997; Moss et al. 1998).

These results on learning feature and conjunction stimuli indicate that differential performance can be evoked

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

with the same set of shapes, according to whether the shapes map onto semantically close or distant representations. It is also possible that contrasts in visual similarity within a set of items also contribute to performance differences, since the semantically close sets used are also very often perceptually similar (e.g., Dixon et al. 1998 showed a deficit for ELM when learning an association between conjunction stimuli and labels for different makes of car – nonliving things within a perceptually close sub-ordinate category). However, it remains an open question whether data derived from a small set of shapes, and tasks that require new learning rather than the retrieval of previously learned information, can be generalised to account for deficits across the broader class of living things. It may be that, within this broad class, perceptual overlap (as well as covariance along multiple dimensions) combines with semantic similarity to create particular problems in identification.

3.4. Visual processing deficits without loss of stored knowledge

As we will document in sections 4 and 6, many patients with category-specific deficits seem to have impairments of stored knowledge, with impaired access to associative and functional (sect. 4) or perceptual knowledge (sect. 6) for the affected items. However, category-specific deficits have also been reported in patients for whom there is no apparent loss of stored knowledge. For example, HJA is an agnosic patient with severe difficulty in recognising many visually presented objects. The deficit is worse for living than for nonliving things, even when these items are matched for name frequency, familiarity, and visual complexity (Riddoch & Humphreys 1987b; Riddoch et al. 1999). Nevertheless, when initially tested HJA showed good stored knowledge for living things, being able to provide detailed definitions and drawings from memory. The only clear problem in stored knowledge occurred when he was asked to retrieve colour information, though this deficit occurred for both living and nonliving things. In his case, it is difficult to attribute the category specific recognition deficit to loss of stored perceptual or functional features; rather it appears to reflect a problem in differentiating between items that have many close perceptual neighbours (i.e., living things relative to nonliving things). This hypothesis fits with the variety of other visual perceptual impairments in this patient (Riddoch & Humphreys 1987b).

Interestingly, when retested some 16 years after the stroke that generated the recognition problems, HJA showed some deterioration in his stored knowledge for the perceptual properties of objects. His drawings from memory were more difficult for independent observers to identify and his definitions listed fewer visual attributes; the decrease in the number of visual attributes listed was more marked for living than for nonliving things (Riddoch et al. 1999). However, there was no general decrease in HJA's performance; he produced more non-visual attributes in his definitions when retested. These data suggest that on-line visual processes interact with memory processes to maintain stored representations of the visual features of objects. Over the longer term, a visual processing deficit leads to some degradation of these memory representations. This may affect representations of living things more because representations of visual attributes need to remain finely tuned to differentiate between these (perceptually similar) items and/or because visual attributes are strongly weighted in our stored representations of living things (sect. 2).

Changes due to learning over time may also explain discrepancies in the performance of one of the first patients reported with category-specific deficits by Warrington and Shallice (1984). In the original study, patient JBR did not show a consistent pattern of deficit across modalities for items probed with pictures and with words, though he was consistently impaired on particular items within each modality. This result has implications for the issue of modality-specific representation of stored knowledge, which we return to in section 6.6. Funnell and de Mornay Davies (1996), however, found that JBR did show consistency across modalities when re-tested some years later. Here learning may have helped him re-establish links for certain items, though there was not a general learning effect as he remained consistently impaired on some items across modalities.

Several other studies, though, cast doubt on the *necessary* relationship between impaired visual and/or structural processing and category-specific deficits. For example, Humphreys and Rumiati (1998) reported data from a patient with suspected Alzheimer's disease, whose problem seemed to reside in poor perceptual knowledge for objects. Though this patient performed well at perceptual match tasks, she was impaired at object decision and on tasks requiring that objects be matched associatively. She also made visual naming errors when objects were misidentified. Despite this, there was no evidence of a category-specific deficit between living and nonliving things. Lambon Ralph et al. (1998) also documented a patient with poor ability to match objects to definitions stressing visual as opposed to functional properties of objects; in addition there was impoverished production of perceptual features in her own definitions. Lambon Ralph et al. found trends for these problems with perceptual information to be worse for nonliving than for living things. Naming also tended to be worse for artefacts. This patient did show good performance on object decision, though, suggesting that there was a problem in interrogating perceptual knowledge from verbal input rather than there being an impairment of perceptual knowledge per se. Hart and Gordon (1992) too reported a case with poor retrieval of perceptual knowledge from names, along with good access from vision, though here the problem was more pronounced for living things.

The contrast between patients who show a deficit for living things due to impaired perceptual/structural knowledge, and those who do not, may be due to differences in the information the patients can draw upon to support performance. The patient of Humphreys and Rumiati, for instance, showed a priming effect on object naming when presented with multiple stimuli from the same category. This is consistent with partial activation of associative knowledge which could boost impaired visual identification when stimuli are presented in an appropriate context. Humphreys and Rumiati proposed that structural descriptions were activated below threshold level unless activation was increased top-down from members of the same category. For patients able to use such top-down activation, living things may benefit owing to their inter-correlated common features; for some patients there may even be relatively strong category activation when single objects are presented, reducing any advantage for nonliving things. Could this also lead to a reverse effect, with nonliving items being the more difficult? Lambon Ralph et al. in fact re-

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

ported a perceptual deficit along with poor matching and production of definitions for nonliving things. In evaluating this, though, it should be borne in mind that, in studies measuring control performance, perceptual definitions of nonliving things are often harder to identify than perceptual definitions of living things (see Humphreys et al. 1997). This may be expected if perceptual information is weighted less strongly than functional (e.g., action-related) information in the stored representations of such stimuli. A patient may perform relatively worse on tests of this kind with nonliving things, then, owing to a general rather than a category-specific decrease in their performance. A test such as drawing from memory, where detailed perceptual information must be retrieved, may enable an impairment with living things to be detected (as indeed was reported by Hart & Gordon 1992).

One other possibility, not tested to date, is that differences between the recognition of living and nonliving things could reflect the ease with which patients can use particular perceptual representations. For example, the agnosic patient HJA has a perceptual deficit which still allows him to derive global shape descriptions, but these representations are not elaborated with local perceptual detail (e.g., see Boucart & Humphreys 1992; Humphreys et al. 1985). Global shape information overlaps across many living things, so that a patient using unelaborated global representations may tend to find these objects especially difficult to identify. Other patients, however, may extract more local detail about the parts of objects and not about the global shape. The parts of nonliving things can have a functional role in their own right (e.g., a wheel of a bicycle), and a patient using local details with such stimuli may sometimes identify such parts as if they are the whole object: a problem we have observed in patients with simultanagnosia. This tendency to identify parts as wholes may be reduced for living things, because either their parts are not functional in their own right or they are diagnostic of the whole (e.g., an animal's head). Patients who "weight" parts more than wholes may find identification more difficult for nonliving than for living things.

As a final point, we note that the prediction, that impaired structural knowledge should *necessarily* disrupt living things more than nonliving things, turns out not to be straightforward – even not taking into account the distinction between global and local object coding. In simulations of the HIT model, we have found that the form of lesioning itself is critical for predicting whether a category-specific deficit emerges for living things after damage to "early" stages of the model. We return to this point in section 6.

3.5. Conclusions

The research reported in this section highlights the importance of visual factors, such as perceptual overlap and covariance in multiple visual dimensions, for at least some aspects of category specific identification impairments. These visual factors may interact with stored representations of visual features (e.g., overlap within a structural description system) and with the semantic retrieval process, to make identification difficult for living things. According to the Cascade model, the combined visual and semantic effects can lead to identification deficits when lesions affect different levels in the object recognition system (access to structural descriptions, access to associative and functional knowledge, and access to object names; Humphreys et al. 1995). Visual deficits seem sufficient to produce "category specific" impairments, but, at the same time, they may not be necessary.

4. Category-specific knowledge

In sections 2 and 3, we have discussed evidence suggesting that category-specific deficits reflect factors other than the categorical nature of our stored knowledge (such as the importance of visual or functional information for representing different categories, or differences in perceptual and semantic overlap between categories). However, the most straightforward account of such deficits is that they arise because our stored knowledge is differentiated according to category of object. We have reviewed Sartori and Job's (1988) claim that stored visual/perceptual knowledge within the structural description system may be categorically organised, and, although we concluded that there was (currently) little empirical support for this idea, we have not discussed the possibility that other types of knowledge are categorically organised. According to the Cascade model, visual/perceptual knowledge in the structural description system is functionally independent of other types of semantic information (e.g., functional information or inter-object associations) and it is possible that this nonperceptual knowledge is categorically organised. We now review this hypothesis in more detail, asking whether this proposal is necessary to account for at least some forms of category-specific disorders.

4.1. Category-specific impairments for nonliving things

One argument for semantic knowledge being categorically organised is that some patients can have deficits for nonliving things rather than living things. Nielsen (1946) first reported this in a patient, CHC, who was able to recognise living things, such as faces and flowers, but not nonliving things, such as a car, a hat, or a telephone. More recently, Warrington and McCarthy (1983; 1987; 1994) outlined three case studies of patients who had particular problems recognizing nonliving things. These patients performed significantly better with living things compared to nonliving things on matching to sample tasks. For example, in spoken word/picture matching, patient VER was better with flowers (93%) and animals (86%) compared to nonliving objects (63%) (Warrington & McCarthy 1983); patient YOT was better with animals (86%) and flowers (86%) compared to nonliving objects (67%) (Warrington & McCarthy 1987); patient DRS was better with animals (95%) compared to nonliving objects (74%) (Warrington & McCarthy 1994). Sacchett and Humphreys (1992) further demonstrated that category-specific impairments for nonliving things can remain when confounding variables such as frequency and familiarity are controlled. Their patient, CW, was significantly better at naming line drawings and at performing picture-word matching tasks for living things (despite the fact that they were less familiar, more visually complex and matched on frequency to the nonliving things). Quite similar data were reported by Hillis and Caramazza (1991), who showed a double dissociation between naming living and nonliving things using the same set of items with two patients. Silveri

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

et al. (1997) also report data from a patient impaired at naming nonliving things using the same stimuli as had been used to elicit a deficit for living things in patient Michelangelo, reported by Sartori and Job (1988). Such double dissociations are consistent with a fractionation of semantic knowledge between that required for the identification of living things and that for the identification of nonliving things.

However, other accounts can be offered for the deficits with nonliving things. Warrington and McCarthy (1987), for example, suggested that impairments could occur at a finer-grained level than the living/nonliving dichotomy. The patient documented in their paper, YOT, was particularly poor at spoken word/written word matching for "small manipulable objects" (e.g., office utensils). They suggested that this was consistent with the view that particular forms of knowledge are needed for (or are differentially "weighted" for) the identification of different objects (see sect. 2). For instance, small manipulable objects might be adversely affected if these items depend on the retrieval of motor memories for their identification and if brain regions supporting these memories were damaged. To account for finer-grained deficits of this sort, we may need to make more than a dichotomous distinction between sensory and "functional" knowledge (cf. Farah & McClelland 1991; Warrington & Shallice 1984). Different forms of both sensory and functional knowledge need to be separated. Furthermore, the "functional" knowledge used in the recognition of living things (e.g., as actors) may differ from that used in the recognition of nonliving things (as instruments).

Now, although it seems intuitively plausible that functional information plays an important role in the identification of nonliving things, there has been surprisingly little behavioural data on the topic (though see sect. 5.1 for correlatory evidence from functional imaging studies). Some direct empirical support comes from Humphreys and Riddoch (1999). They examined a child, JS, with learning difficulties due to a genetic abnormality and birth trauma. JS was presented with pictures of artefacts along with either a contextual scene or with an action (mimed by the experimenter). In the scene condition JS was told that "this is an 'n' and you find it in an 'm'" (e.g., "this is a glass and you find it on a table," where the glass was shown alongside a pictured table with a place setting). In the action condition she was told "this is an 'n' and you do this with it" ("this" being shown by the mime). JS was asked to repeat the object's name, and then the next learning trial began. Subsequently she was presented with the individual objects in isolation and asked to recall the names. Performance was better after learning with the actions than the contexts. The finding is of some interest because it suggests that pairing an artefact with an action is useful in establishing a stable link between the object and its name. Interesting: JS was better at naming nonliving things relative to living things. Humphreys and Riddoch proposed that nonliving things might benefit from having object-specific action pairings, which can help name retrieval. In contrast, the actions performed by living things will tend to be more similar and hence less useful in supporting naming of the individual stimulus. Note also that the actions performed *by* living things are also different in kind from the actions performed *on* nonliving things. This also may be important for learning in such a child.

Most studies of patients with deficits for nonliving things have not clearly defined the stage in the recognition and naming process at which the deficit might occur, using tasks designed to "tap" structural, functional/associative (semantic) or name knowledge. The patient reported by Cappa et al. (1998), however, showed no effect of category when answering probe questions, and they suggest that the problem was in name retrieval for nonliving things. Other patients, though, have shown a deficit in discriminating between nonliving targets and semantic distractors (e.g., Sacchett & Humphreys 1992), which is more suggestive of a deficit in accessing associative knowledge. It may be, as with deficits for living things, that we will need to account for impairments arising at different levels of the object processing system.

4.2. Impaired functional knowledge for living things

Accounts that stress the importance of forms of functional knowledge only for nonliving things have some difficulty with the finding that some patients show impaired retrieval of functional as well as visual/perceptual information about living things. In contrast, the same patients may demonstrate reasonably good functional knowledge about nonliving things (Caramazza & Shelton 1998; Funnell & de Mornay Davies 1996; Laiacona et al. 1993; Samson et al. 1998; Sheridan & Humphreys 1993). If functional knowledge is differentially important for identifying artefacts, then loss of this knowledge ought to lead to poor identification of these items. Clearly this is not *necessarily* the case. The data can be accommodated, though, by the idea of categorically organised stored knowledge. For example, Caramazza and Shelton interpreted the above pattern of data in terms of there being separate semantic representations for living and nonliving things, so that damage to the representations for living things leads to poor retrieval of both stored perceptual and functional knowledge for these stimuli. Caramazza and Shelton (1998) suggested that this separation of knowledge for living and nonliving things reflects evolutionary pressures that highlight the importance of the categories animal, plant life, and artefact. They argued that, because of evolutionary gains in distinguishing between these three types of object, categories for animals, plants, and artefacts form the basis for the organisation of conceptual knowledge. "The evolutionary adaptations for recognising animals and plant life would provide the skeletal neural structures around which to organise the rich perceptual, conceptual, and linguistic knowledge modern humans have of these categories."

This view, that stored knowledge is organised by these three basic categories, can account for some of the finer-grained dissociations found between patients showing category specific deficits. For instance, whilst some patients seem particularly poor with fruits and vegetables (Farah & Wallace 1992; Forde et al. 1997; Hart et al. 1985), others have been reported with deficits only for animals and not for fruit and vegetables (Caramazza & Shelton 1998; Hart & Gordon 1992). This would be expected if knowledge about animals and plant life is functionally (and anatomically) separate. For accounts that do not assume that stored knowledge is categorically organised, such finer-grained impairments could be attributed to the use of particular forms of knowledge to identify particular objects; perhaps shape information is more important for animals and colour or texture for fruit and vegetables. Loss of information

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

about shape or colour knowledge might selectively impair recognition of animals or fruit/vegetables, respectively (see sect. 4.1 above).

Consistent with Caramazza and Shelton's hypothesis, Keil (1987) has argued that the underlying conceptual structures for living and nonliving things are qualitatively different, and will be affected in qualitatively different ways by the transformations that can be applied to them. Keil used the example of a chair that has its back sawn off and an extra leg glued on – most people would agree that it has become a stool rather than a chair. However, dyeing a dog's fur red, making its tail bushier, and allowing it to live in the wild, will not change it from a dog into a fox. By changing the salient perceptual and action-related features of nonliving things we can alter identities; the same does not apply for living things. When asked to justify why, for example, the dog did not become a fox, participants stated that it was because the internal organs (or DNA) would remain the same (although they could not state how the internal organs of foxes and dogs differ). Keil proposed that "rather it is a belief in biological essence that seems to grow out of a naïve theory of natural kinds that is driving their intuitions."

The argument for knowledge being organised along categorical lines has been given further support by studies of semantic development in children.[3] S. Gelman (1988) reported that four-year-olds can explicitly state whether objects are made by people or not (see also S. Gelman & Kremer 1991), and a number of studies have demonstrated that preschoolers (again, around four years old) have a relatively sophisticated understanding of the differences between living and nonliving things. For example, R. Gelman and Meck (R. Gelman 1990) asked children to describe what was on the inside and outside of animate and inanimate objects. The insides and outsides of animate things were described in different ways, whilst the insides and outsides of inanimate objects were described in the same way. R. Gelman (1990) argued that these results highlight children's tendencies to generalise properties across examplars of animate objects, but not across the animate-inanimate boundary, and proposed that knowledge about animate and inanimate objects was domain-specific and governed by contrasting underlying rules.

A similar argument comes from the work of Massey and R. Gelman (1988). They presented children with colored photographs of novel examples of mammals, nonmammalian animals, rigid complex artefacts, wheeled objects and statues that had animal-like parts, and asked them which items were capable of going up a hill by themselves. Three- and four-year-olds chose the animals, and rejected the artifacts with animal-like parts. This indicates an ability to use the appropriate visual-perceptual properties present in the novel objects to decide whether or not the object was capable of self-initiated movement. R. Gelman (1990) argued that this ability to distinguish between objects that can and cannot move on their own is the basis for the development of the conceptual distinction between animate and inanimate objects.

However, a number of other studies have indicated that the animate/inanimate distinction is acquired even earlier in infancy (Mandler & McDonough 1996; Poulin-Dubois et al. 1995; Smith 1989). Smith (1989) presented 12-month-old infants with pictures of nonobjects that were made from combining parts of animals (to make "living nonobjects") or artefacts (to make "nonliving nonobjects"), and, using a habituation paradigm, demonstrated that the children could distinguish between the "living" and "nonliving" things. Mandler et al. (1991) further showed that 18-month-old infants could make distinctions at an even more fine-grained level, separating animals, plants, and nonliving things. Indeed, even within the domain of nonliving things the infants could apparently distinguish furniture from kitchen utensils, though they could not separate tools and musical instruments. These studies suggest that infants are able to make judgments about global categories very early, perhaps even before they can categorise objects at a basic level (e.g., a chair as opposed to furniture). Mandler et al. (1991) posited that infants do not use physical similarity when forming categories but rely instead on more abstract properties, such as movement patterns and origin of movement (self generated or extrinsic), which are associated with each category (see also R. Gelman 1990; Mandler 1992). Consistent with this, Mandler and McDonough (1996) found that 14-month-old infants generalised their responses to actions between living things, but not to nonliving things (and vice versa). When shown a dog drinking, the infants were willing to make a rabbit drink but not a motorcycle. When shown a car being started with a key they were willing to start a truck with a key but not a fish.

These findings suggest that even young infants can distinguish between different categories, lending support to the idea that there is a predisposition to differentiate between living and nonliving things. Furthermore, if living and nonliving things do come to belong to different knowledge domains, perhaps represented in different neural areas, then we would expect category-specific patterns of impairment after brain damage.

4.3. Some problems

The proposal that our semantic knowledge is categorically organised has primarily been driven by reports of patients who have category specific impairments accessing visual and functional information about living things, with no concomitant problems with nonliving things. However, these patients are only problematic for accounts stressing the importance of visual/perceptual knowledge for living things, and functional knowledge for nonliving things, if the functional knowledge evaluated is the same for both categories. But, as we have already noted, the term "functional" knowledge has been used in neuropsychological studies to refer to all kinds of nonsensory information and does not always refer strictly to the function (or use) of the object. Typically "functional" knowledge for nonliving things refers to object usage and motor activity on the part of the actor; in contrast, "functional" knowledge for living things includes biological functions (e.g., eating), the context in which animals are found, the sounds animals make, and so forth. Hence the "functional" knowledge that can be impaired in patients with impairments for living things is not the same as the "functional" knowledge about nonliving things that can be spared. Until there have been reports of patients with category specific problems in accessing visual and functional knowledge when the type of information is matched across categories, there is no empirical reason to abandon an account which stresses the importance of different types of information for different categories. Furthermore, we suggest that accounts stressing the importance of different

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

types of knowledge for different categories can in fact account for patients who have category specific impairments accessing *all* kinds of semantic knowledge (i.e., "functional" as well as visual/perceptual). Let us suppose that living things are represented primarily in terms of visual and sensory features. Now, if these important (perhaps even defining) visual/sensory features are impaired, then patients may not have enough information to differentiate one exemplar from another. For example, if a patient does not know that a giraffe has a long neck, he does not really know what a giraffe is, and therefore could not answer a question tapping "functional" knowledge, such as "Does a giraffe eat meat or leaves?" This hypothesis has been supported by recent neuroimaging studies, which show that modality-specific areas associated with the processing of form are activated when subjects answer verbally presented questions about the visual and *functional/categorical* properties of living things (Chao et al. 1999; Thompson-Schill et al. 1999; see sect. 5). It appears that perceptual knowledge of form is fundamentally important for our stored representations of living things. "Knowing" what a living thing is, and being able to answer all kinds of questions about it, depends heavily on being able to access this stored perceptual information. In contrast, activation in cortical areas associated with form processing is not enhanced when subjects answer questions about the functional/categorical properties of nonliving things. This is consistent with the idea that the important contrasts, or defining attributes, are visual/perceptual for living things but not necessarily for nonliving things. Chao et al. (1999) in fact report increased activation in the middle temporal gyrus for tools (relative to animals) when categorical information must be retrieved. They point out that this region borders areas known to be specialised for processing motion (area V5), suggesting that information about motion when tools are used plays an important role in their categorisation.

We also suggest that the developmental research is consistent with the idea that different types of knowledge are important for different categories. Indeed, R. Gelman (1990) argued that children learn to differentiate between animals and non-animals on the basis of the different movement patterns shown by these categories. Furthermore, the categorical account of stored knowledge does not distinguish between different types of learned representation (visual/perceptual and semantic/functional, etc.), and consequently, has difficulty explaining why visual processing disturbances can be linked to poor identification and loss of stored knowledge for living things (sect. 3). Also, the account has difficulties in accommodating cases where functional knowledge for living things appears to be intact and the deficit is only for the visual properties of these objects (Hart & Gordon 1992; Riddoch & Humphreys 1993). In section 6, we outline the HIT account, that allows these different lines of evidence to be integrated.

5. Anatomical considerations

5.1. Brain imaging studies

A number of recent studies have attempted to assess whether living and nonliving things are stored in anatomically separate brain regions, using functional imaging techniques. Martin et al. (1996) used positron emission topography (PET) to compare the brain regions that were active when subjects named living (animals) and nonliving (tools) things. When they subtracted the regions activated when subjects were naming tools from the areas activated when subjects were naming animals, they found that naming animals selectively activated the left medial occipital lobe and inferior temporal regions. When they reversed the subtraction (i.e., areas active when naming tools minus areas active when naming animals) they found that naming tools selectively activated the left middle temporal gyrus and the left premotor region. Interestingly, the areas selectively activated when subjects were naming tools were very similar to those activated when subjects named actions associated with objects (Martin et al. 1995) or when they imagined grasping objects (Decety et al. 1994). This is consistent with the idea that activating information about object use and the associated motor action is important for naming tools. In contrast, the activation of stored visual information, within the left medial occipital and inferior temporal cortex, may be more prolonged and important for the identification of living things. Again, this idea is supported by converging PET data indicating that similar areas are activated when subjects retrieve stored colours associated with objects (Martin et al. 1995). The results support the notion that different brain regions mediate the identification of living and nonliving things. Furthermore, since converging evidence can indicate the functional roles of the different brain regions, the results highlight the importance of particular domains of knowledge for certain objects (respectively visual and action-related, functional knowledge for living and nonliving things).

Quite similar results to those reported by Martin et al. (1996) were found by Perani et al. (1995), using a task in which pairs of objects had to be matched according to whether they had the same base-level names (e.g., a tennet and a hacksaw). Damasio et al. (1996), like Martin et al. (1996), examined functional activation in an object naming task, but focused on activation within regions of the temporal lobe. They found that naming animals resulted in activation of the left inferotemporal region (more anterior to regions highlighted by Martin et al. 1996) whilst the naming of tools led to activation of posterior middle and inferior temporal gyri.

It is not clear whether some of the reported anatomical differences reflect variations in scanning procedures or in the task requirements; nevertheless, the data highlight the importance of the inferior temporal lobe for naming animals and the left, inferior frontal and posterior middle temporal regions for naming tools. There may, in addition, be selective involvement of the left medial occipital area in animal identification.

One problem with these PET studies is that the stimuli have tended not to be matched across the categories. Differences in the activations found with animals and tools, then, may be due to the contrasting types of knowledge invoked by these stimuli or differences in stimulus complexity or familiarity. Stronger activation of posterior visual processing areas (e.g., in medial occipital cortex) may be found with animals because they are visually more complex. Moore and Price (1999) attempted to assess this by comparing PET activations in naming and picture-word matching tasks with visually complex animals, visually simpler fruit/vegetables, visually complex (multi-component) nonliving things (vehicles, appliances) and visually simpler (single-component) nonliving things (tools, utensils). They

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

found increased activation in the right medial extra-striate and occipito-temporal areas for multi-component relative to single-component items, and particularly for animals relative to the other categories they examined. This suggests that these posterior brain regions are implicated in the processing of more complex visual configurations, with animals being particularly dependent on activation here. In addition, living things (animals and fruits and vegetables) were associated with activation of the left anterior temporal cortex and nonliving things with activation of the left posterior middle temporal cortex (see also Damasio et al. 1996; Martin et al. 1996). These results suggest that, instead of thinking of there being one area specialised for processing living things, and one for nonliving things, there are rather multiple areas, implicated in processing and retrieving different forms of knowledge for the particular stimuli. For example, for living things, areas supporting both visual processing and the retrieval of associative knowledge may be important. Thus extra-striate areas are sensitive to visual complexity, and they are also implicated in the processing of animals, consistent with there being a strong mediating role of visual knowledge in the identification of these stimuli. Activation of the anterior temporal cortex for living things, similar to that reported by Damasio et al. (1996), may instead link visual processing to the retrieval of stored associative knowledge. This fits with lesion data from a patient (reported by Breedin et al. 1994) who suffered profound atrophy of the anterior, inferior temporal cortices whilst leaving more posterior areas of temporal cortex intact. This patient had only a mild deficit on object decision tests designed to assess access to stored structural knowledge, and he showed strong perceptual priming (Srinivas et al. 1997), along with poor judgments based on semantic relatedness between objects. His performance was also worse with living than nonliving things. For accounts such as the Cascade model (sect. 3.1), living things should generate stronger associative as well as perceptual competition between stimuli, requiring prolonged differentiation at both perceptual and associative levels, as indicated by different sites of activation in functional imaging studies and effects of brain lesions at contrasting sites.

To link data on functional imaging to models in a more precise way, we need to establish the brain sites implicated at different stages of object processing. Price et al. (1996), for example, assessed whether there were particular brain regions linked to object naming as opposed to object recognition (referring here to access to stored structural knowledge). They had subjects perform four tasks: (1) name (coloured) objects, (2) name the colour of nonobjects (matched to the object for complexity), (3) say "yes" to the coloured objects (the object baseline), and (4) say "yes" to the coloured nonobjects (the nonobject baseline). They assumed that objects gain automatic access to structural information (cf. Glaser 1994), and argued that the contrast between the two baseline conditions (with objects and nonobjects) can reveal areas involved in object recognition (activated by objects but not by equally complex nonobjects). The brain regions mediating name retrieval may be indicated by the contrasts between the two naming conditions (conditions 1 and 2) relative to the two baselines (because both the naming tasks required retrieval of a stored phonological label associated with the visual stimulus). Of most relevance here, though, is the interaction given by the contrast between the object naming condition and its baseline (conditions 1 minus 3) relative to the contrast between the nonobject colour naming condition and its baseline (conditions 2 minus 4). This interaction indicates areas associated specifically with retrieving the names of known objects from vision. Price et al. found that the interaction was linked to selective activation of the left inferior and posterior temporal lobe. This study suggests that this region is particularly involved in name retrieval for known objects, and it is more activated in name retrieval than in object recognition alone. The same area has also been found to be activated more by animals than by tools (Martin et al. 1996; Perani et al. 1995). If such inferior and posterior temporal regions are linked to the involvement of stored visual knowledge in processing, the data suggest that this knowledge is recruited particularly when name retrieval is involved. The conclusion that follows from this, even in online tasks, is that knowledge sources may be activated to different degrees, in a top-down manner (e.g., for naming, but not for recognition). We return to this in our formulation of the HIT account.

Studies using fMRI have added extra information by differentiating further between regions within the posterior temporal cortex and by contrasting activation from objects with that found when the same knowledge is addressed from verbal questions. Thompson-Schill et al. (1999) proposed that stored visual knowledge was not only activated when subjects were asked to name living things, and to answer questions concerning their visual attributes, but also when asked to answer questions about non-visual information (e.g., Are pandas found in China? Are snails edible?). In contrast, they argued that it was not necessary to access stored visual information when functional questions were asked about nonliving things. In their study, subjects were asked yes-no questions about the visual and functional properties of living and nonliving things, and fMRI was used to assess whether the left ventral occipito-temporal cortex (particularly the fusiform gyrus) was activated in each condition, relative to a baseline condition in which nonsense auditory stimuli were presented. The left ventral occipito-temporal cortex was chosen as the area of interest because previous studies have shown that it is involved in representing the visual properties of objects (D'Esposito et al. 1997; Martin et al. 1995). Thompson-Schill et al. found that there was a significant interaction between category (living vs. nonliving) and type of question (visual vs. functional) in the pattern of activity in the fusiform gyrus (assumed to be involved in representing modality-specific visual/perceptual information). In particular, they found that, while this area was only activated by the visual questions for nonliving things, it was activated by *both* the visual and functional questions for living things. Chao et al. (1999) likewise report a remarkable overlap of areas in the lateral fusiform gyrus activated by animals both in an object naming task and in a task requiring both naming and the retrieval of meaning from words (e.g., to answer the question, "forest animal?" To the word "deer"). They found similar overlap across tasks for tools, but in different areas – the middle fusiform gyrus and middle temporal gyrus.[4] These results point to common regions being activated from different modalities. The regions implicated in the study of Chao et al. have, in other experiments, been linked to the processing of form (in the lateral fusiform gyrus) and the processing of motion (in middle fusiform and temporal ar-

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

eas) (see Bonda et al. 1996; Ungerleider & Haxby 1994). Hence the findings are consistent with specific forms of knowledge being drawn on to different degrees when accessing stored knowledge about living and nonliving things, and that these forms of knowledge are organised respectively around sensory and motor processes. It also fits with this argument that areas of the left inferior frontal cortex, when scanned, have been shown to be activated more by nonliving than living things (and typically the nonliving things have been tools). From other imaging studies it appears that left inferior frontal cortex is implicated in the retrieval of knowledge about actions, and it is activated when subjects must name actions associated with tools as well as when they name the tools themselves (Grabowski et al. 1998; Grafton et al. 1997; Martin et al. 1995). Chao et al. add that additional differences are apparent within the categories of living and nonliving items (e.g., between houses, tools, and chairs), reflecting the kinds of knowledge important in the representation of each type of stimulus (e.g., houses, though nonliving, are not associated with activation of the middle temporal gyrus). We return to the implications of these data in section 6.

In summary, neuroimaging data indicate that living and nonliving things activate different areas of cortex in identification tasks. These different areas of cortex can also be linked to contrasting stages of identification: visual processing and access to visual knowledge (bilateral, medial extra-striate), contrasting with access to associative (anterior temporal), motion, and action knowledge (in left medial temporal and fronto-parietal regions). Naming, as opposed to recognition, also seems to involve activation of additional visual processes, consistent with there being top-down recruitment of visual knowledge for individuation of a target object from its perceptual neighbours (particularly in left inferior occipito-temporal regions).

These conclusions are also consistent with evidence on category differences in event-related potentials (ERPs). ERP studies are particularly useful for providing information about the time course over which representations are activated. An ERP study by Kiefer (in press) suggests that living and nonliving stimuli generate both contrasting and overlapping ERPs, at different times. Subjects had to judge whether a super-ordinate category probe (presented verbally or by means of two pictures from the chosen category) was appropriate to a target stimulus (presented either as a name or a picture). Kiefer found that early ERP differences emerged for living and nonliving targets (after about 160–200 msec), which were restricted to when pictorial presentations were used. These early effects are consistent with there being more elaborated perceptual processing of pictures of living things than nonliving things. There were also differences between the categories on later ERPs, with living and nonliving things respectively producing a reduced N400 component over occipito-temporal (bilateral) and fronto-central (left hemisphere) sites. Other ERP studies have demonstrated differences in fronto-central sites for N400 components associated with action verbs versus concrete nouns (Dehaene 1995; Pulvermuller et al. 1996), suggesting that the fronto-central changes are linked to the retrieval of associated action knowledge. These differences in the later ERP component in Kiefer's study were found with both words and pictures as stimuli. The late effects are consistent with living and nonliving things drawing on contrasting forms of associative knowledge, stored in different brain regions, for both verbal and pictorial stimuli alike.

5.2. Lesion studies

In addition to their PET data, Damasio et al. (1996) reported results from group studies of patients with deficits in naming living or nonliving things (see also Tranel et al. 1997). They found that impairments in naming animals were associated with damage of the left inferotemporal region, and impairments in naming tools were associated with damage to the junction of the temporal, occipital, and parietal cortices. Damasio et al. (1996) argued that their data reflected naming rather than recognition problems, since patients were classified according to whether independent raters could identify the object from the description produced by patients when the object could not be specifically named. In contrast, Tranel et al. (1997) reported patients classified by recognition impairments (after unilateral right as well as left hemisphere damage). They reported that recognition problems for tools were associated with damage to the left occipito-temporal-parietal junction, whilst recognition problems for animals were associated with both left and right medial occipito-temporal lesions. The lesion sites linked to recognition problems for animals in this last study were more posterior to those linked to naming problems in the earlier study, though conclusions need to be cautious. Correct identification of object descriptions by independent raters (as in Damasio et al. 1996) does not necessarily indicate that the patients had intact access to semantic knowledge. Also in these studies patients were classed as showing deficits if they performed at a level of 2 standard deviations below that shown by control subjects; however, this does not necessarily mean that the patients were more impaired on one category of object than another, when the categories were compared directly (see Caramazza & Shelton 1998). Nevertheless, we note that the recognition deficits associated with posterior right as well as left hemisphere medial occipito-temporal damage link to Moore and Price's (1999) finding that multi-component objects (and particularly animals) strongly activate this brain region. It may be that additional visual processing is needed in this area to differentiate within sets of complex, structurally similar objects, generating both enhanced activation in normal subjects and deficits in recognizing animals in brain lesioned patients. This also fits with evidence on the category-specific deficits for living things found in the patients reported by Forde et al. (1997), Humphreys et al. (1997), and Riddoch et al. (1999). These patients had posterior occipito-temporal damage and all had particular problems in interrogating visual knowledge about objects. For instance, the patients were poor at object decision, at answering definitions that stressed the visual properties of objects, and at drawing objects from memory. These problems were more pronounced for living things than for nonliving things. Associative and functional (action related) knowledge about both living and nonliving things was good.

More anterior temporal regions may be involved in access to associative knowledge, necessary for name retrieval for living things (see also Breedin et al. 1994). Consistent with this, deficits in some patients with poor identification of living things is linked to bilateral damage to anteromedial parts of the temporal lobes as well as to more inferior temporal regions (Gainotti et al. 1995). However, the

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

degree to which the damage *needs* to be bilateral is unclear. Caramazza and Shelton (1998) for example report a patient with impaired knowledge of living things following unilateral damage confined to the left hemisphere (though to fronto-parietal regions in this case). The patients reported by Forde et al. (1997) and Humphreys et al. (1997) also suffered unilateral left hemisphere lesions (but, as noted above, to relatively medial occipito-temporal areas). For these patients, the deficits for living things were apparent in naming rather than recognition tests (e.g., fruits and vegetables could be categorised but not named). This suggests that left hemisphere representations are implicated particularly in naming tasks, with posterior occipito-temporal regions being interrogated to differentiate targets within sets of living things.

In contrast to the data on deficits with living things, Gainotti et al. (1995) conclude that deficits in recognising and naming nonliving things occur after damage to left parietal-frontal regions. However the data base for these patients is somewhat limited. Silveri et al. (1997) report one such case where an analysis using PET was conducted. They showed hypometabolism confined to the middle temporal gyrus, the hippocampal and parahippocampal regions, and the inferior parietal lobe of the left hemisphere. In contrast, Cappa et al.'s (1998) patient, who apparently had problems particularly in naming nonliving things, had a unilateral lesion affecting the anterior left temporal lobe. Tippett et al. (1996) similarly reported problems in naming nonliving relative to living things in patients with resections of the left anterior temporal lobe to relieve epilepsy.

Overall, the neuropsychological data concur with the results from functional imaging studies in suggesting that different brain regions mediate the recognition and naming of living and nonliving things. Lesions affecting inferior occipito-temporal regions, extending anteriorly into the temporal lobe, seem particularly to disrupt processing for living things and lesions affecting left temporal-parietal and also parietal-frontal areas disrupt processing for nonliving things. There are some anomalies though, such as the left fronto-parietal lesion linked to a deficit in living things reported by Caramazza and Shelton (1998), and the data suggesting that the anterior left temporal area may be involved in naming nonliving things (Cappa et al. 1998; Tippett et al. 1996). These inconsistencies illustrate the importance of collecting further data, across larger sets of patients each assessed to provide a functional analysis of any deficit relative to a model of object naming.

6. The hierarchical interactive theory (HIT)

The data we have reviewed have demonstrated the reality of category-specific deficits in object naming as empirical phenomena, with effects occurring for both living and nonliving things. The majority of accounts that have been outlined to explain these deficits have emphasized single factors – the categorical nature of our stored knowledge (sect. 4), the differential roles of perceptual versus functional knowledge for living and nonliving things, respectively (sect. 2), the importance of correlated perceptual features or of perceptual-functional relations (sect. 2), the interaction of perceptual processes with other forms of knowledge needed for a task (sect. 3). Our interpretation of the evidence is that, whilst particular deficits can be linked to each account (so that each is *sufficient* to explain particular patients), no single factor account can cover all of the patients (i.e., no one account is *necessary*). Instead, we propose a framework that is powerful enough to accommodate different patterns of deficit in contrasting patients.

6.1. Hierarchical knowledge and interactive processing

We begin by discussing the need for a hierarchical system for object processing by outlining a model of visual object recognition and naming, as this is the task most commonly used to diagnose category specific impairments. We then discuss the need for top-down (interactive) processes within such a model. These two ideas, of a hierarchical and interactive system, capture the essence of the HIT account. To account for object naming, HIT adopts the architecture of the Cascade model of object recognition, which had two defining characteristics: a hierarchy of stored representations and the assumption that partial activation could be transmitted between processing systems. The hierarchy of stored representations in the Cascade model contained three types of stored knowledge: (1) stored structural descriptions, (2) stored functional and inter-object associative information ("semantic" knowledge), and (3) name representations. According to the model, selective damage can occur to each form of representation, so that, for example, a patient can have a pronounced deficit in semantic knowledge (e.g., poor retrieval of associative or functional knowledge about an object) without necessarily having a deficit in stored structural knowledge. As we outline below, this is consistent with dissociations between the performance of different patients on tasks designed to "tap" each stage of memory storage. Note, however, that although a hierarchy of memory storage is assumed, processing within models of this type does not depend on access to each memory stage being completed before the next is initiated. Thus, as far as processing is concerned, only a first pass of activation is hierarchical. Processing, after this first pass, is not necessarily hierarchical. For example, there may be a delay in accessing precise structural knowledge from an object, though access to categorical knowledge may be completed (see Humphreys et al. 1997).

Let us consider the neuropsychological evidence for a hierarchy of memory storage. Take first the distinction between stored structural descriptions and other types of knowledge. Evidence for this comes from patients who can perform object decision, but who still show impaired access to semantic information about an object. In object decision tasks patients can be asked to discriminate between stimuli (of equal perceptual complexity and perceptual "goodness"), some of which are real objects and others nonobjects created by interchanging the parts of real objects. Because objects and nonobjects cannot be distinguished on the basis of general perceptual information, their classification must depend on association with stored knowledge. Some patients can succeed at object decision whilst being impaired at matching associatively related objects or accessing any further information about the entity (Hillis & Caramazza 1995; Humphreys & Riddoch 1999; Riddoch & Humphreys 1987c; Sheridan & Humphreys 1993; see also Breedin et al. 1994). This suggests that such patients can access structural knowledge (to distinguish objects and nonobjects) even when access to associative knowledge is deficient.

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

Some patients with category specific deficits for living things are poor at performing object decisions, especially for items from the affected categories (e.g., Caramazza & Shelton 1998; Sartori & Job 1988). However, this is not universally the case, and some patients demonstrate intact object decision performance for affected categories (Laiacona et al. 1997; Sheridan & Humphreys 1993). We propose that these contrasting patients have lesions affecting different functional stages of object processing. The former "category specific" patients have deficits affecting visual access to structural knowledge (in the structural description system) in addition to any further deficits affecting functional/associative knowledge or name retrieval. The latter "category specific" patients do not have a deficit affecting the visual activation of structural knowledge, and their deficit resides at a later processing stage.

There are other explanatory advantages in a model separating stored visual/perceptual knowledge, in the structural description system, from other types of knowledge. In particular, such a model can account for patients who show good knowledge of functional/associative knowledge of living things when given their spoken name, whilst still being impaired in naming them from vision (Forde et al. 1997; Riddoch & Humphreys 1987c; 1993). Damage confined to the structural description system should leave other forms of knowledge intact, though visual naming (and access to visual knowledge in general) may be impaired. Such a model can also help explain a pattern of deficit in which a patient has good visual access to structural knowledge (e.g., object decision), good auditory access to functional/associative knowledge, but poor auditory access to structural knowledge (e.g., drawing from memory). This pattern was reported by Hart and Gordon (1992) and Riddoch and Humphreys (1987c). It can be accommodated if there is a selective impairment in mapping back from the semantic system to the structural description system, and if the structural description system is used when long-term knowledge of visual properties must be retrieved.

All of the above deficits can be conceptualised in terms of hierarchical memory representations that are accessed in a feed-forward manner, as proposed by the Cascade model (Humphreys et al. 1988) – providing there is also off-line interrogation of visual/perceptual knowledge in the structural description system. The latter process would be required when we are asked to retrieve perceptual knowledge when given the name of an object. We next consider a pattern of performance that suggests that, even for on-line naming, some form of re-interrogation of structural knowledge is required. This leads to our formulation of the HIT model.

As we discussed in sections 3.1 and 5.2, some authors have reported patients whose category specific problems seem more pronounced on name retrieval rather than on tests of semantic knowledge (Cappa et al. 1998; Farah & Wallace 1992; Forde et al. 1997; Hart et al. 1985; Humphreys et al. 1997). Nevertheless, in a detailed evaluation that has been conducted on two such patients with problems with living things, the fundamental problem appeared to be in activating perceptual knowledge about objects. As we have noted, Humphreys et al. (1997) reported two patients who were impaired at naming living things but, on forced-choice tests, showed apparently good access to associative and categorical knowledge about these objects (e.g., successfully categorising fruits and vegetables). Yet, despite this good performance on associative and categorical tasks, deficits in perceptual knowledge were revealed by tests of drawing from memory, object decision, and naming to perceptual definitions. The apparent inconsistency between, on the one hand, the relatively good access to associative and categorical information from vision and, on the other, the perceptual knowledge impairment, may be accounted for in at least two ways. It may reflect the ability of patients to access partial associative and categorical knowledge and to respond accurately using this information under forced-choice conditions. In addition, it may reflect the role of top-down (re-entrant) activation of perceptual knowledge in object naming, which we discuss below.

The idea here is that, in a first pass in object processing, there is activation of stored structural descriptions and partial activation of associative/functional knowledge. However, for naming to be achieved there needs to be further (top-down) interrogation of perceptual knowledge, which drives the process of differentiating a target object from its close neighbours. For living things, this may require further interrogation of form information for animals, colour and texture for fruits and vegetables, and so forth. For nonliving things this may mean further interrogation of action-related functional knowledge, that distinguishes one artefact from another. In patients such as those described by Humphreys et al. (1997), with problems in naming living things, we propose the following. There is initial visual access to associative/functional knowledge, but the mild impairment in perceptual knowledge prevents successful re-entrant activation from being achieved. Consequently, patients cannot access enough information to differentiate between the target item and its close perceptual and semantic neighbours. The HIT model incorporates this idea of re-entrant processing, which can be conceptualized in terms of an interactive activation and competition framework, as illustrated in Figure 2. This distinguishes the model from a pure feed-forward account, such as the Cascade model of Humphreys et al. (1988). According to HIT, there is not only activation of a hierarchy of processing stages in object naming, but also interactivity in this processing so that different forms of knowledge are activated in a reiterative fashion for name retrieval to operate.

We can link this idea of re-entrant processing to the PET results on object naming reported in section 5.1. There we discussed the data of Price et al. (1996), showing increased activation of the inferior, posterior left temporal lobe in object naming relative to object recognition. We suggest that this increased activation reflects the extra visual processing needed for name retrieval, as opposed to object recognition, to take place.

We have attempted to explore the reasons why top-down re-entrant processing may be useful in object naming through simulations. As previously noted, Humphreys et al. (1995) implemented a simple interactive activation and competition model of object naming, which incorporated top-down as well as bottom-up connections between each level within a hierarchy of memory representations (for structural, "semantic" and name information about objects). The model captures the hierachical and interactive nature of HIT, though its knowledge representations are undoubtedly oversimplified and primitive; nevertheless it allows us to undertake a preliminary exploration of the dynamics of processing different categories of object within this framework. Objects from living and nonliving cate-

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

gories were represented in the model by vectors reflecting the rated perceptual similarity between objects. Thus for any given target object, its own structural representation would be assigned a value of 1 whilst structural representations of other similar objects would be activated too, but to a lesser degree (see sect. 3.1). Naming was assumed to occur when units at the name level reached a set activation threshold (+0.2). Overlap between the perceptual descriptions of living things slowed their naming, because it created competition for access to specific name representations (relative to when nonliving things were presented). Within such a model, living things can be named either by (1) prolonged perceptual processing (enabling the visual input to over-ride any initial competition at the name level) and/or (2) re-entrant "clean-up" which reduces overlap within the structural description system. Watson and Humphreys (in preparation) explored the utility of re-entrant clean-up in this model first by selectively adding noise to top-down connections, and second, by removing top-down connections from the semantic to the structural representations (so that the model operated solely in a feed-forward fashion, from structural descriptions to semantics). For the HIT approach to be viable, the top-down connections ought to have some functional effect on object naming. Apart from the alterations to the strength of top-down connections (through noise and selective disconnection), the parameters of the model were otherwise the same as those reported in Humphreys et al. (1995)[5]. Consistent with the HIT approach, Watson and Humphreys found that re-entrant clean-up was generally useful for speeding up the naming process, and particularly so for living things. Varying degrees of noise were added to the weights connecting "semantic" to structural units (i.e., to disrupt re-entrant activation to the "structural description system"). When the average noise was 0.25 per connection (for a maximum connection strength of 1), "RTs" for living things increased from a mean of 24.2 to 34.1 cycles; "RTs" for nonliving things remained constant (20.6, taking an activation value of 0.2 on the output units as a threshold). Increasing the top-down noise further (e.g., to a mean parameter value of 5 per connection) led to errors being made (where name units failed to reach the set threshold after 1,000 cycles). The identification rate for living things fell to 50% whilst that for nonliving things remained at 100%. In this case prolonged processing of stimuli was insufficient to enable identification to take place, when top-down (or re-entrant) noise was present. Similarly, removing the connections from semantic to structural units also disrupted the naming of living things, though less dramatically (RTs for living things increased from 24.2 to 30.6 cycles; for nonliving things there was no change). These simulations indicate that, where stimuli overlap perceptually, re-entrant processing can be useful in differentiating the input efficiently, for individual identification to take place.

6.2. Category specific deficits

We have argued that a model with hierarchically-arranged memory representations is required in order to account for the selective impairments in retrieving particular forms of knowledge in different patients (e.g., to account for patients who can access perceptual but not associative or functional knowledge normally from vision). We have also proposed that re-entrant top-down processing is useful for distinguishing between perceptual neighbours, when object naming is required. We next review how a framework assuming interactive processing, within a hierarchical set of memory representations can accommodate the category specific nature of the resultant deficits.

First, any model that incorporates re-entrant feedback for on-line naming may suppose that activation is transmitted continuously between processing levels. From this follows the same arguments concerning category specific deficits for living things made by the Cascade model (Humphreys et al. 1988). Living things suffer due to perceptual overlap between representations of exemplars from the same category. This generates competition between perceptual representations, and also between representations that are accessed "down-line" (e.g., for associative and functional knowledge, and for names), all of which will slow name retrieval. This competition is present within the system in its "normal" (unlesioned) state, and it can be exacerbated by damage to various levels of the naming system (see Humphreys et al. 1995, for simulations). On top of this, there may be additional, detrimental effects for living things when stored perceptual knowledge is degraded, if there are indeed correlated relations (and hence strong inter-connections) between their perceptual and functional features (McRae et al. 1997; Moss et al. 1997; see sect. 2). With such inter-correlated representations, damage to some features may generate impairments across a range of stimuli. For nonliving things, similar damage should be less effective because features will be common to smaller sets of items.

In addition to effects of perceptual overlap and feature correlation, perceptual features tend to be more diagnostic for living relative to nonliving things (for which action-related functional knowledge may be particularly diagnostic). The shape of an animal or the colour of a fruit will carry more information about the identity of that object than, say, the precise shape or colour of a kettle. Due to these factors, we suppose that visual information will have a higher weighting than other types of information in our representations of living things. Hence damage to such features should be particularly disturbing. For an analogous reason, damage to action-related knowledge may selectively impair nonliving things.

Finally, the HIT model supposes that knowledge is re-activated in order to facilitate name retrieval. We suggest that re-activation of perceptual knowledge is particularly important for the naming of living things for at least three reasons: (1) owing to their greater perceptual overlap, re-entrant "clean up" processes are more likely to be useful (Watson & Humphreys, in preparation); (2) such processes are more necessary because of correlated perceptual features and/or correlated perceptual and functional features, (3) perceptual features are more diagnostic (and more strongly weighted) than functional features for the identification of living things. Deficits in re-entrant activation of perceptual knowledge can account for the selective naming problems, without apparent recognition problems, for living things in at least two patients reported in the literature (Humphreys et al. 1997).

To date, our discussion has highlighted re-entrant activation of visual knowledge, which we suggest involves brain regions in the inferior occipito-temporal regions. However, whilst this may be useful for identifying living things (owing to their perceptual overlap, and the importance of dis-

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

tinctive perceptual features for defining these items), re-entrant activation of different forms of stored knowledge may be useful for identifying other objects. As has been noted by other authors (Warrington & McCarthy 1987), sensori-motor knowledge (e.g., concerning actions that might be associated with the objects) may be useful for differentiating many nonliving things. For this reason, re-entrant activation of brain regions representing this knowledge (e.g., at the temporal-occipital-parietal junction, and in inferior left frontal regions, Damasio et al. 1996; Grabowski et al. 1998; Martin et al. 1996) may be involved in identifying nonliving things. As a consequence of this, deficits in naming nonliving things could be caused by damage to the neural areas that are the sites of this re-entrant activation. Note that a mild impairment at this level may allow some semantic knowledge to be retrieved even if naming fails (Humphreys et al. 1997). This may account for patients of the type reported by Cappa et al. (1998) and Silveri et al. (1997), who performed well on word-picture matching tasks whilst being selectively poor at naming nonliving things. A more severe impairment would prevent even associative/functional knowledge from being retrieved, as well as causing a naming impairment. Overall, however, and in keeping with the hierarchical nature of our account, we suppose that more posterior damage will be linked with problems in accessing and/or re-activating visually-related knowledge (though the visual information involved may differ across categories; e.g., form for animals, colour/texture for fruits/vegetables, motion for tools, etc.)(see Breedin et al. 1994).

There are however at least two problems for the argument that re-activation of visual/perceptual knowledge for form and colour is particularly influential for recognising and naming living things. These are that: (1) some patients with impairments for living things have deficits with functional as well as visual/perceptual knowledge, even when assessed auditorily (Caramazza & Shelton 1998; Laiacona et al. 1993; 1997; Sheridan & Humphreys 1993); and (2) some patients with impaired visual/perceptual knowledge do not have a category-specific deficit for living things (Humphreys & Rumiati 1998; Lambon Ralph et al. 1998). We deal with each in turn.

Concerning the first point, the HIT account suggests that visual/perceptual knowledge will be drawn upon even when answering functional and associative questions about living things, to the extent that visual/perceptual knowledge is strongly weighted within the overall, distributed representation of such stimuli. This assertion is supported by data from functional imaging (Chao et al. 1999; Thompson-Schill et al. 1999). Hence it is possible for a deficit of visual/perceptual knowledge to influence the retrieval of non-sensory knowledge. In addition, though, the model holds that forms of associative and functional knowledge may be represented separately for different categories of object, not least because this knowledge varies across categories – we note the earlier point we made concerning how functional knowledge for living and nonliving things can differ in kind. This means that associative and functional knowledge may be selectively impaired for living things even if visual/perceptual knowledge is intact. The patient reported by Sheridan and Humphreys (1993) illustrates this. She had poor functional knowledge about fruits and vegetables yet discriminated real fruits and vegetables from perceptually similar nonobjects in a difficult object decision task. This indicates a deficit in functional but not visual/perceptual knowledge, which is still specific to living things. Thus non-sensory knowledge about living things can be disrupted by damage at either a visual/perceptual level (which impairs access to semantic knowledge) or at a functional level. The extent to which a visual/perceptual deficit impinges on the retrieval of functional knowledge will also vary across patients. A mild visual/perceptual deficit may leave the retrieval of functional information relatively intact, as documented by Forde et al. (1997; see also Humphreys et al. 1997).

The second point, concerning patients with visual/perceptual impairments that are not category-specific, is more difficult, since at first sight the HIT account predicts that visual/perceptual damage will differentially influence living things. In section 3.4, we discussed some ways in which theories could explain how visual/perceptual deficits might affect living and nonliving things to an equal degree, and the same points could be reiterated here (e.g., where patients are able to use semantic knowledge to recover information for living things, or where different types of perceptual information are damaged). In addition to this, our simulations with HIT have indicated that the exact predictions depend on the type of damage that is inflicted (Watson & Humphreys, in preparation). As we noted in section 6.1, when we add noise to the connections to and from the structural description system to functional and associative knowledge, the identification of living things becomes increasingly (and selectively) problematic. However, if "dynamic" noise is added, which fluctuates over time as activation is cycled through the system, then a different pattern of results emerges. First, the detrimental effects of the noise are smaller. Second, the differential impairment to living things is decreased. For example, when a maximum noise value of 5 per connection was added on random iterations, the mean identification time for living things increased from 24.2 to respectively 33 cycles (for bottom-up noise, from structural descriptions to associative knowledge) and 32.7 cycles (for top-down noise, from associative knowledge to the structural descriptions). The mean identification time for nonliving things increased too, though, from 20.6 to 26 iterations (for both bottom-up and top-down noise). With dynamic noise, the network is never pushed toward a single value in vector space but is rather biased randomly in different directions over time, as the network converges. This slows identification time, but generally to an equal degree across stimuli, irrespective of whether stimuli come from a part of vector space where representations are "sparse" (for nonliving things) or "clustered" (as case for living things). "Static" (constant) noise, however, pushes activation in the network systematically away from the target representation. There is then greater disruption for objects represented in terms of similar vector values (for living rather than nonliving things). The simulations suggest that we need to have precise ways of linking lesions to models, in order to formulate predictions. For now we simply note that HIT does not *necessarily* hold that damage to visual/perceptual knowledge selectively impairs living relative to nonliving things.

6.5. Implementing HIT and its relation to other models

HIT bears a clear family resemblance to several other models in the literature. Like the Cascade model (Humphreys

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

et al. 1988) it allows activation to be transmitted continuously between processing levels and it emphasizes the knock-on effects of perceptual and semantic overlap on the recognition and naming of objects. It is also consistent with ideas that correlated activation of perceptual features, and strong links between common perceptual features and functional information, can play a determining role in recognition and naming (cf. McRae et al. 1997; Moss et al. 1997). In this respect, the model captures many of the ideas expressed in the OUCH model of semantic memory (Caramazza et al. 1990), which holds that perceptual-functional associations lead to rapid access to semantic knowledge for objects. HIT diverges from the latter models in specifying a hierarchical architecture of memory representations, separating stored structural descriptions from semantic knowledge to better account for fractionations between patients (sect. 6.1). HIT also diverges from all of these accounts in the proposal that re-entrant activation plays an important role, particularly in name retrieval. The idea that re-entrant activation plays a role in object identification is similar to the arguments made previously by Warrington and McCarthy (1987). They suggested that different modalities of experience (visual, tactile, auditory, kinesthetic, etc.) contribute in retrieving the meaning of objects, and that objects can vary in how they are represented in terms of their weighting across the different modalities. Visual/perceptual knowledge may be weighted more strongly than functional knowledge for living things; functional knowledge may be more strongly weighted for nonliving things. Damasio and colleagues (Damasio 1989; 1990; Damasio et al. 1996) have similarly argued for distributed forms of object representations, involving limbic as well as modality-specific forms of knowledge (limbic representation of an objects' affective association, visual representation of its appearance, auditory representation of any associated sound, verbal representation of abstract knowledge, etc.). These distributed knowledge sources will be activated in a variety of tasks, depending on the extent to which an object must be differentiated from competitors for the task to be achieved. Evidence on the facilitatory role of action on the learning of names to artefacts (Humphreys & Riddoch 1999) provides evidence for a modulatory role of action on name retrieval for these stimuli.

To date, our simulations of HIT have used an interactive activation and competition framework (Humphreys et al. 1995; Watson & Humphreys, in preparation). This has the benefits of having a modular representation scheme, in which representations at different processing stages can be damaged selectively (e.g., so that access to associative/functional knowledge can be damaged but not access to visual/perceptual knowledge). However it also has some disadvantages. For instance, the localist representation scheme used makes it more difficult than some approaches to incorporate robust forms of learning, though it is possible to use forms of Hebbian learning to establish weightings based on correlated activation between representations (see Burton 1994; also Pulvermuller [1999] makes similar arguments for a role of Hebbian learning in establishing learned representations for objects across contrasting neural regions). Other models have approached simulations of category-specific deficits using fully distributed representations and generalised learning routines such as back propagation (e.g., Durrant-Peatfield et al. 1997), or using self-organising feature maps, which can establish local "fields" for particular object categories and so can be lesioned selectively (e.g., Mikkulainen 1997; Ritter & Kohonen 1989; Zorzi et al. 1999). Within the last approach, some of the problems in identifying living things can be understood in terms of a coarsening of the process of mapping from perceptual descriptions of objects into maps representing associative and functional knowledge. When stimuli overlap in terms of both their perceptual and associative/functional representations, performance will be particularly disrupted (see Dixon et al. 1997, for a similar argument, couched in terms of Kruschke's [1992] ALCOVE model of object categorisation). An equivalent approach can also be adopted on the output side, for mapping from associative/functional knowledge to names. Damage to these output maps should result in a deficit that is more pronounced on naming than on recognition (see Zorzi et al. 1999). Such an "output" deficit could occur in addition to output problems stemming from impaired re-entrant activation of forms of knowledge critical for identification, and should thus be associated with different lesion sites (e.g., for living things, this output deficit may be linked to more anterior lesions than a deficit in re-entrant activation; cf. Forde et al. 1997; Humphreys et al. 1997, who report apparent output deficits along with posterior lesions).

This discussion makes it clear that the HIT approach could be implemented in a number of ways, using a variety of architectures; indeed Figures 3 and 4 here illustrate different ways in which stored knowledge could be represented within the framework. Critical to the approach, though, are two points: (1) that there is a hierarchy of processes, so that (e.g.) perceptual knowledge can be spared when other forms of knowledge are damaged, and vice versa; and (2) that re-entrant activation is useful for identification. These two critical points enable us to explain: (a) those patients with intact access to perceptual but not to other forms of stored knowledge (Hillis & Caramazza 1995; Humphreys & Riddoch 1999; Riddoch & Humphreys 1987c; Sheridan & Humphreys 1993), (b) apparent naming impairments for living things following inferior occipito-temporal lesions, (c) the greater activation of inferior occipito-temporal regions when living things are identified, both from vision and audition, and the association of these areas with category-specific deficits for living things in patients, and (d) the greater activation of left middle fusiform/temporal gyrus and inferior frontal areas when nonliving things are identified (particularly tools), and the association of these areas with category-specific deficits for nonliving things.

Now, given that HIT could be implemented in a number of ways, it could be argued that the framework is in fact too powerful, and not open to refutation. Against this, the above two points do stress critical issues on which the approach could be refuted: if no fractionations occur between access to contrasting knowledge sources in different patients (e.g., access to perceptual but not functional knowledge), and if naming could not be linked in a relatively direct way to the recruitment of visual knowledge over and above that found in categorization tasks. However, on these points, we suggest that the evidence is supportive.

6.6. Processing in other modalities

Our discussion of HIT has focussed on the naming of visually presented objects, since this is the task that has been used in most studies in the literature. However, there seems

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

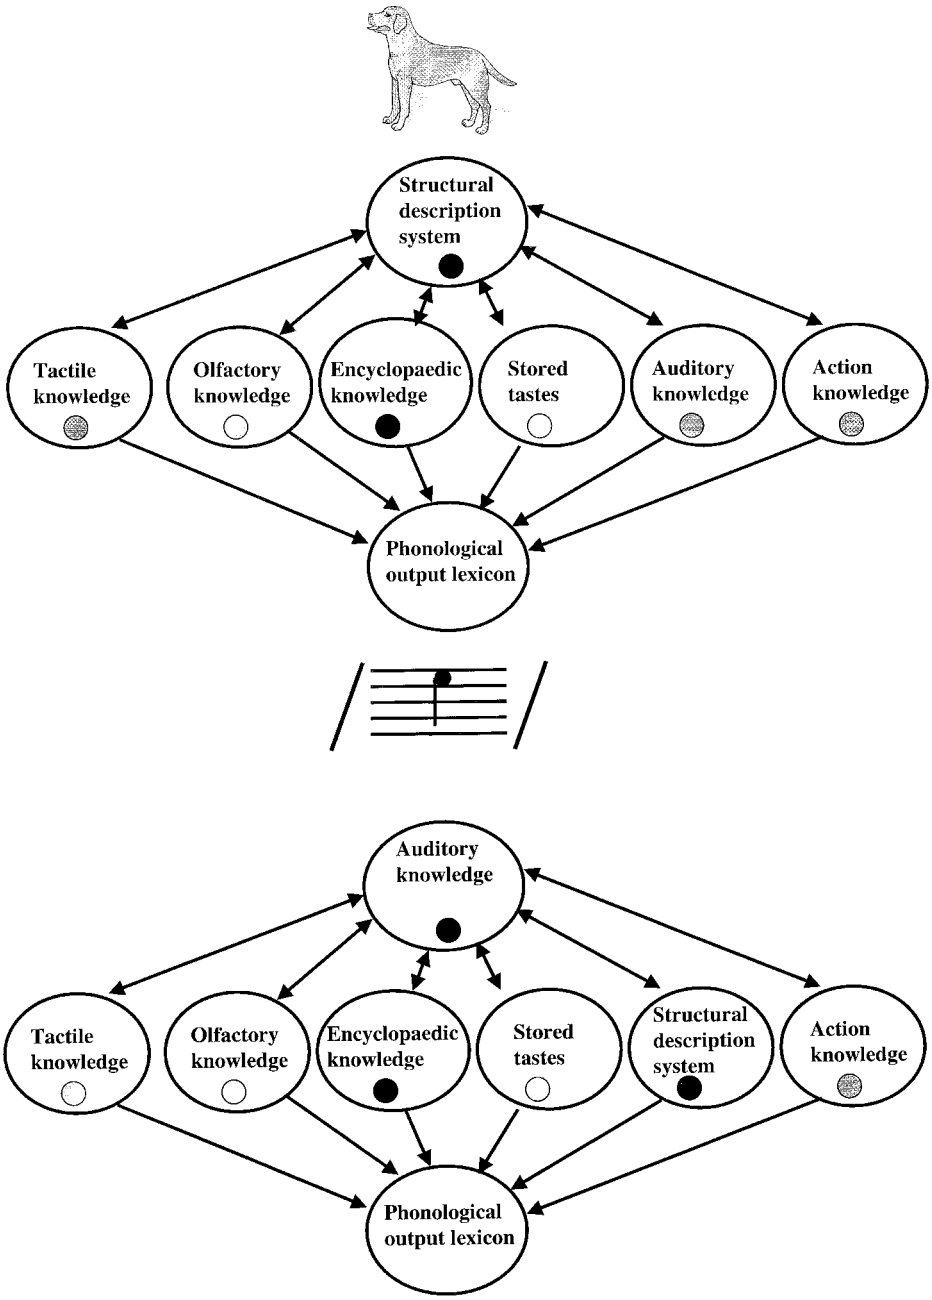


Figure 3. (Top) A framework for HIT illustrating the semantic representation of an animal (e.g., a dog), accessed from vision. Within this scheme there are distinct forms of stored knowledge for various forms of perceptual, motoric, and encyclopedic knowledge about objects. The dog is represented in terms of a pattern of activation across these different forms of stored description, shown here by the colouring given to a "unit" within each knowledge store (the darker the unit, the more strongly activated). We posit stored perceptual descriptions for action, for sound (auditory knowledge), for taste, smell, and touch. Note that the units do not necessarily represent localised knowledge nodes, but could represent a pattern of distributed activity across multiple units within a particular subsystem. Phonological representations could then be activated on the basis of a given pattern of activation across all areas.

(Bottom) A framework for HIT illustrating the semantic representation of a musical instrument (e.g., a guitar), accessed by means of its sound. Note that both the type of information accessed first, and the pattern of distribution across the different types of semantic knowledge, differ from the example in Figure 3(Top).

no reason why a similar architecture could not be adopted to account for processing in other modalities. Perceptual recognition stores (equivalent to the structural description system) may exist for touch, smell, taste, sounds and so on (as illustrated in Fig. 3). These different perceptual knowledge stores would be interrogated to retrieve different forms of stored knowledge about objects. Other nonperceptual forms of knowledge may exist in a further verbal knowledge store (e.g., for encyclopedic knowledge such as who sang "Old Shep"). Naming can be contingent on the appropriate activation pattern being established across the different knowledge stores. Note that, on this account, all forms of knowledge would not necessarily be accessed when naming takes place (see Chertkow et al. 1992, for some evidence). Furthermore, on this account these modality-specific forms of perceptual knowledge would not only be

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

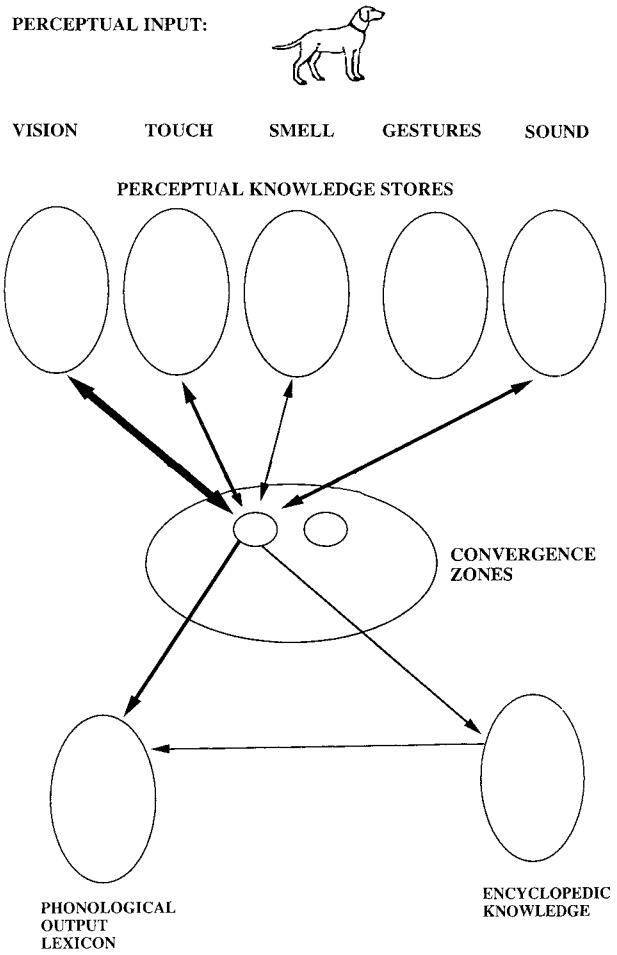


Figure 4. An alternative framework for HIT where "semantic memory" is captured in terms of convergence zones. The model here illustrates the semantic representations of a dog. The amount of information stored in particular sensory areas, and the importance of this information, will vary across different categories of object (as indicated by the thickness of the lines).

contacted on-line in identification tasks (within each respective modality), but they would also be accessed off-line when input from another modality must draw upon specific forms of knowledge. For example, to decide whether an elephant has a relatively long or short tail, when given its name, we must access stored visual knowledge. To decide if a lemon is sour or sweet, when given its name, one might access the same sensory knowledge that allows us to recognise a lemon when we taste it. For different categories of object, we believe that the activation of certain perceptual knowledge sources is more critical than others.

Now, if both on-line and off-line tasks involve activation of diverse knowledge systems (some perceptual, some sensori-motor, etc.), we may come to question whether the idea of "semantic memory" becomes superseded. What is semantic memory other than the sum of its distributed perceptual, motor, verbal, and connotative components (see Allport 1985; Warrington & McCarthy 1987, for similar views)? In discussing object naming we were careful to distinguish structural from functional and associative knowledge, and Figure 3 suggests that similar distinctions can be made for other input modalities. But, in what sense is functional knowledge more "semantic" than the structural knowledge that defines the form of an object – especially if we define some forms of functional knowledge in terms of the action associated with an object (as with tools)? For living things in particular, connections from the structural description system to other knowledge systems may indeed be central in our long-term representations, and more important than other types of knowledge for defining these objects. Does the idea of a "semantic memory," perhaps organised by biological category, have any reality over and above this?

We suggest two approaches to this question. One is to maintain that certain knowledge stores (e.g., a verbal store specifying encyclopedic knowledge, not gained by direct sensory experience with stimuli) are truly "semantic." These knowledge stores may indeed represent something like a dictionary definition of the meaning of objects, which is accessed irrespective of the modality of the input. Another stems from the arguments of Damasio (1989; 1990). Damasio proposed that distributed knowledge stores (perceptual, sensori-motor, verbal, etc.) are bound together through "convergence zones," which store the combinatorial arrangements of the different forms of stored knowledge for each object. Object naming involves reciprocal activation between modality-specific knowledge systems and the convergence zones, which serves to synchronize neuronal firing and stabilize memory retrieval (see Fig. 4). We may even conceive of these convergence zones in terms of topographic maps that link different forms of sensory and motor representations (see sect. 6.5). One interesting aspect of this proposal is that communication via convergence zones may be necessary to access inter-object associative knowledge even within a modality (e.g., the structural descriptions for cup and saucer may be linked via convergence zones). It follows that damage to this process will selectively impair access to inter-object associative knowledge even when access to stored perceptual knowledge of an individual object is intact (see Hillis & Caramazza 1995; Riddoch & Humphreys 1987c). In the framework as set out in Figure 3, it is not clear how inter-object associations within a modality are represented.

It is also possible that convergence zones and/or verbal encyclopedic knowledge stores are categorically organised (Caramazza & Shelton 1998), which would allow for both "true" category specific deficits and for deficits that mimic category specificity whilst reflecting a co-varying factor (e.g., impaired differentiation of visual features). However, we have argued that there is currently little empirical evidence of "true" category specific disorders. Alternatively, convergence zones that represent exemplars from the same category could be functionally/anatomically close because they link to similar sensory areas (Damasio 1990). Convergence zones for strawberries, apples, and peaches will have interactive links with cortical areas representing information about smell, taste, colour, and shape whereas the important links from the convergence zones for saws, hammers, and screwdrivers will be to areas storing the associated motor engrams, functional information, and shape.[6] On this view, convergence zones moderate access to forms of information other than those that can be derived directly from the modality-specific sensory input. Damage across convergence zones will lead to a close coupling of impairments across modalities (e.g., Forde & Humphreys 1995; 1997), whilst damage to knowledge systems supporting access to a given zone produces modality-specific disorders.

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

Note that damage to these zones cannot be thought of as losing stored representations per se, since the zones serve as access devices rather than knowledge stores. In accessing knowledge systems for particular tasks, contrasting neural areas will be drawn upon, as suggested by the functional imaging data (sect. 6). The degree of activation of a given neural area (visual auditory, sensori-motor, etc.) will reflect factors such as the degree of differentiation required along that dimension (e.g., the degree of perceptual overlap) and the weighting of that dimension in the knowledge representation (these last factors being closely related).

6.6. Conclusions

We have argued that (1) there is no single locus of category-specific impairments for either living things or nonliving things; several functional loci can each be identified as *sufficient* to cause the deficits, but none seems necessary; and (2) stored knowledge is represented in a distributed manner, with different forms of perceptual and action-related knowledge being recruited on-line for particular tasks. At the most, "semantic" memory may reduce to our verbal, encyclopedic knowledge about objects; minimally it may be no more than the currently instantiated patterns of activation across the distributed knowledge representations. To accommodate the findings, we have proposed the HIT. It brings together a hierarchical model of memory representation with an interactive approach to knowledge retrieval, for object naming. The interactive approach to processing enables different forms of knowledge to be used for particular objects and tasks, and it can be used to accommodate neuropsychological evidence on lesion sites and functional neuroimaging data on the involvement of contrasting brain areas in the naming of different classes of object.

ACKNOWLEDGMENTS

This work was supported by grants from the Medical Research Council, the European Union, and the Human Frontier Science program. The ideas have benefited from discussion with many people but we especially thank Jane Riddoch, Cathy Price, and Derrick Watson.

NOTES

1. We acknowledge that an even greater range of deficits than this can be distinguished in the neuropsychological literature, including deficits with abstract relative to concrete nouns, with colour names, action names, proper names, and so forth (see Cohen & Burke 1993; McCarthy 1995). However, the literature on selective deficits with living and nonliving objects is itself both large and diverse, and requires extensive discussion in its own right. Hence we confine ourselves to this topic. It is possible that these other deficits can be conceptualised within the frameworks we discuss.

2. The term structural description here is meant to primarily to distinguish a perceptual memory of the visual form of objects from associative, contextual, and functional (e.g., action-based) knowledge of objects (typically thought to be represented in semantic memory). We are not committed to whether this perceptual memory explicitly represents parts of objects and parts-based relations, or whether it is more holistic in nature.

3. Note that these studies have focused on how children learn to distinguish between animate and inanimate objects rather than the more general living-nonliving dichotomy.

4. Chao et al. contrast activation across categories but not within categories relative to a baseline condition, such as passive viewing. Hence it is difficult to assess whether visual areas, such as the lateral fusiform gyrus, were activated relative to baseline when answering non-visual questions about nonliving things (though relative to the activation produced by nonliving things, answering non-visual questions about living things *did* further activate this region). Thompson-Schill et al. (1999) included baseline conditions within subjects, to enable comparisons to be made within each category.

5. Connections form semantic to name units were set as being either "high" (0.9) or "low" (0.7), to represent objects with high or low frequency names, respectively.

6. Note that in our model nonliving things (such as tools) will have links to visual/perceptual areas (e.g., representing shape). The point is that visual/perceptual information will not be as important for these items because (1) there may be fewer links (e.g., no representation of colour, fewer parts) and (2) visual/perceptual attributes alone may be less useful for differentiating a target object from its competitors.

Open Peer Commentary

Commentary submitted by the qualified professional readership of this journal will be considered for publication in a later issue as Continuing Commentary on this article. Integrative overviews and syntheses are especially encouraged.

Shortcomings of the HIT framework and possible solutions

Martin Arguin and Daniel Saumier

Groupe de Recherche en Neuropsychologie Expérimentale, Département de psychologie, Université de Montréal, Montréal, Canada and Centre de Recherche, Institut Universitaire de Gériatrie de Montréal, Montréal, Canada.
arguinm@psy.umontreal.ca **saumierd@hotmail.com**

Abstract: Although the principles defining HIT find support within a variety of fields, the framework is currently underspecified in several respects, which prevents it from fully achieving its purpose. We highlight a number of such shortcomings and propose avenues for a resolution of some of these issues.

Category specific deficits in brain damaged individuals point to the interactive effects of the organization of our cognitive system and of the structure of the world we inhabit on object processing performance. Humphreys & Forde's (H&F) review describes several factors that may determine this interaction. The HIT framework they propose is defined by few very general principles that are supposed to allow the expression of each of these factors, in the hope of producing a general account of category specific deficits.

The simulation studies reported by H&F demonstrate that HIT may account for deficits specific to living things that result either from a visual encoding impairment, loss of structural knowledge, or damaged connections between the structural description system and subsequent processing stages. In every case, the deficits rest on the assumed greater visual similarity between semantically related living things than that between nonliving things (e.g., Arguin et al. 1996; Dixon et al. 1997; Humphreys et al. 1988). The effect of similarity may plausibly be extended to other types of features, which could differentially affect processing across categories and modalities.

Similarity alone appears insufficient to account for the complete variety of category specific deficits reported so far. This account seems to require, in addition, an appeal to other factors,

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

such as the variable weights different types of knowledge may have for the representation of particular categories (e.g., Farah & McClelland 1991; Warrington & Shallice 1984), and the pattern of feature correlations that exists within and between representation domains for objects of different categories (e.g., McRae et al. 1997; Moss et al. 1997). H&F do appeal to these factors. They do not however, emerge spontaneously from the current definition of HIT and a more detailed formulation appears necessary to allow their implementation.

The weight of a particular type of knowledge for the representation of objects within a category may be assumed to depend on whether the items from that category register within this domain of knowledge representation. This is probably determined by two factors: (1) whether an object's features are relevant within that domain (e.g., actions for using a fork vs. a giraffe); (2) the stability of features within that domain (e.g., color of lemons vs. telephones). There is presently no provision within HIT to allow these crucial determinants to affect the organization of the knowledge representation system. Some plausible extension to the theory specifying how the experience of feature relevance and stability may affect the information that is stored about an object therefore appears required for HIT to naturally implement the factor of variable weights of knowledge types across categories.

The notion that correlations between object features may affect processing concerns the degree to which the presence of a particular property in an item predicts other properties. Within a neural network such as HIT, this effect would most likely be mediated by the connection weights between feature pairs. A specification of how these variable connection weights may plausibly be implemented (presumably on the basis of learned conditional probabilities) will be necessary before the concept of feature correlations can be considered applicable within the HIT framework.

Other relevant issues are the sequence of processing stages within the hierarchy proposed by HIT and the constraints that may apply on the connections between the assumed representation systems. These are presently poorly specified. These issues however, seem relevant with respect to HIT's account of deficits specific for nonliving things. This account appears to rest on the assumption that the role of action-related functional knowledge in the processing of nonliving things is directly analogous to that of visual structural knowledge that was demonstrated for the recognition of living things. Within the context of the visual object recognition task however, it may be noted that whereas the correspondence between an input image and its appropriate structural description is rather systematic, the correlation between this unprocessed image and the actions we may perform with the depicted object appears extremely weak. This is why theories of visual object recognition generally assume that access to stored visual knowledge is required to mediate the access to more abstract forms of knowledge (Biederman 1987; Tarr & Bülthoff 1998). Even without this assumption though, it is clear that any direct access of action knowledge from the input image cannot be as effective as that of structural knowledge, given the differing degrees of transparency in the mappings involved. A difference between action and structural knowledge in terms of the ease and/or order in which they are accessed from an input image means they must play unequal roles in the visual object recognition task, independent of the category the target object belongs to. How this unequality may then interact with the living versus nonliving distinction is a complex issue that requires a more detailed treatment than that offered by H&F.

Finally, a fundamental assumption of HIT is that knowledge about the world is distributed across multiple stores specialized for specific types of information. H&F present an empirical basis supporting a distinct store for visual structural knowledge. The foundations for the other knowledge stores (e.g., action or biological function knowledge) discussed by H&F are unclear, however. An apprehension of the different kinds of properties that may characterize objects is an insufficient basis to assume multiple segregated knowledge stores. In the absence of empirically demonstrated dissociations, general guiding principles should also be involved in making deductions regarding this issue. For instance, it may be noted that, in a neural network, knowledge is stored mainly within the connection weights, that serve notably to convert the representation of an object from one type of format into another (e.g., from visual structure to meaning). Therefore, a consideration of the functional (e.g., representation formats involved and the transparency of the mappings between them) and neurophysiological (e.g., neural pathway development and synaptogenesis) factors that determine the connections available within the network may be part of the decision criteria for a psychologically plausible type of stored knowledge.

Making living versus nonliving distinctions: Lessons from infants

Martha E. Arterberry

Department of Psychology, Gettysburg College, Gettysburg, PA 17325.
arterber@gettysburg.edu **www.gettysburg.edu/~arterber/sheet.html**

Abstract: Developmental research on infants' categorization of living and nonliving objects finds that very young children are equally skilled in grouping such objects. The lack of a specialization for one type of object over another may be due to knowledge of function and the time frame for acquiring such knowledge.

Investigations of how our representational systems fall apart, as presented by Humphreys and Forde (H&F), provide insights into the nature of intact structures and possibly their development. At first glance, H&F's three-tiered model (HIT) for object knowledge appears to be consistent with the developmental progression of concept acquisition in infancy. It is likely that infants' early concepts are based on perceptual similarities among category members. Some claim that, within the first year, infants' categorization is based on responsiveness to perceptual or structural features (Bornstein 1984; Mandler 2000). For example, facial configurations contain important information for categorizing cats and dogs (Quinn & Eimas 1996). Infant attention to object function and category assignment based on function may emerge in the second year. Infants group animals and vehicles based on the presence of legs or wheels at 14- and 18-months-of-age (Rakison & Butterworth 1998), and by 18 months infants recognize correlations between the form and function of parts of artifacts (Madole & Cohen 1995). Last in development comes facility with naming. Infants' object naming explodes after the acquisition of approximately the first 50 nouns, something that generally occurs between 18- to 20-months-of-age (e.g., Miller 1981). Thus, at least, in this superficial sense we find developmental support for H&F's model.

A further look at the developmental literature reveals that from the start humans' categorization of living and nonliving objects may be on equal footing. Categorization research with the youngest of participants has been conducted with 3-month-olds, and some of this work has assessed their ability to categorize living things in comparison to nonliving things. Three-month-olds categorize mammals and furniture presented as still images (Behl-Chadha 1996) and animals and vehicles in still images and moving point-light displays (Arterberry & Bornstein 2000). An interesting finding in both of these studies was that an advantage for one object type, living or nonliving, was not found over the other – a finding that is in contrast to some other areas of infant perception and naming studies with intact adults. Schiff et al. (1989), for example, found differential performance by 5-month-olds in an intermodal perception task with living (a person) and nonliving (a car) objects, and Laws and Neve (1999) found an advantage for naming living things in "normal" adults.

One reason we do not see a specialization for one category of objects over another in very young infants may be due to how the

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

relation between perception and function is viewed. In the HIT model, visual/perceptual information (stored structural descriptions) is separated from functional (or action) knowledge, and the latter is a more advanced process than the former. The authors claim, as an example, that to name living objects, more attention is paid to form, whereas to name nonliving objects, more attention is paid to action. However both form and action information is available at the visual/perceptual level, and according to some theorists (e.g., Gibson 1979), visual/perceptual information cannot be dissociated from function (or action) information. From this view, "function" refers to what is possible based on the needs of the actor and the properties of the object. Moreover, young infants respond to this information (e.g., Adolph et al. 1993).

By contrast, H&F use the term "function" to refer to conventional use. Adults have had years of experience to learn the conventional uses and actions of objects, and over time they may disregard what is possible in favor of what is typical. To continue with H&F's example in section 6.2, a kettle is typically used for cooking, but it can function as a vessel for hauling dirt or other substances. The property of containment is available in the perceptual array and is immediately available to the observer, regardless of age or experience. However, what substance one particular culture contains in the kettle is learned over time. Thus, acquisition of such knowledge may further refine the organization of the concept.

The role of increasing knowledge in concept organization is illustrated by a recent theory of children's acquisition of the distinction between animacy and inanimacy. Several researchers have suggested that an understanding of animacy and inanimacy in general, and the animal-vehicle distinction in particular, is based on the knowledge of the differences between how animate and inanimate objects move (Gelman 1990; Mandler 2000; Rakison & Poulin-Dubois, in press). Rakison and Poulin-Dubois (in press) posit that over the first two years of life children acquire several levels of understanding as they come to fully appreciate the distinction between animate and inanimate objects. The distinction begins with an early understanding of the onset of motion (self-propelled vs. other caused) and lines of trajectory (smooth vs. irregular) of different classes of moving objects, and it culminates with an understanding of the purpose of action (such as goal-directed vs. without aim) and the influence of mental states (intentional vs. accidental). Thus, with development, increasing knowledge of function and/or the animacy distinction, may lead to the living-nonliving specialization that is found in adulthood.

ACKNOWLEDGMENT

The author is also a Visiting Scientist at the Child and Family Research Section of the National Institute of Child Health and Human Development, Bethesda, Maryland.

Is category specificity in the world or in the mind?

H. Clark Barrett

Center for Adaptive Behavior and Cognition, Max Planck Institute for Human Development, 14195 Berlin, Germany. **barrett@mpib-berlin.mpg.de**

Abstract: HIT produces category-specific deficits without category-specific mechanisms by assuming that differences in properties of objects are transparently converted into differences in representational format. A complete model would specify the mechanisms that accomplish this. Such category-specific mechanisms may have evolved because assumptions about the properties of some kinds of objects (e.g., living things) are invalid for others (e.g., artifacts).

HIT is one of a class of models that postulates category-specific deficits in the absence of category-specific mechanisms. Instead, the model contains a variety of information formats (e.g., perceptual, semantic), none of which is category-specific in itself. In order for category-specific deficits to emerge from such a model, there must be principled differences in the kinds of information used to represent different kinds of object. How do these category-specific differences in representational format arise?

In many models, including HIT, different types of information are assumed to be privileged for different categories of object, a priori (see also Farah & McClelland 1991). For example, Humphreys and Forde (H&F) suggest that living things have high "perceptual overlap" (i.e., high within-category feature correlation, also called "structural similarity"), and that "perceptual features tend to be more diagnostic for living relative to nonliving things . . . The shape of an animal or the colour of a fruit will carry more information about the identity of that object than, say, the precise shape or colour of a kettle. Due to these factors, we suppose that visual information will have a higher weighting than other types of information in our representations of living things" (sect. 6.2).

But where does this "weighting" come from? Because assumptions such as these are built into the model, category-specific deficits appear to emerge for free, as the mechanisms that produce them are not specified. In other words, category-specific differences are assumed to be transparently available properties of the world, rather than of the mind.

Clearly, if there are category-specific differences in how objects are represented in different information stores, damage to a particular store will systematically impact different categories of object. But what produces this category-specific privileging of particular information formats? Is it learned? Are some kinds of information more "available" than others in the stimulus objects themselves? Or might there be evolved mechanisms that privilege certain kinds of information for particular categories of object, precisely because this information is more useful for the purposes of identification, inference, and so on?

Reason to suspect the latter lies in two observations. The first is that capacities to identify and make inferences about particular naturally occurring kinds of object, including taxonomic living kinds, tool-like artifacts, food substances, and so on, would have been adaptive and strongly selected for in the past. Second, it is precisely because there are kind-specific differences in principles that can be used to diagnose category membership, and to make adaptive inferences about category members (including those H&F have pointed out), that we expect mechanisms for constructing specialized representations of particular object kinds to have evolved. A brief examination of the principled differences between biological taxa and artifacts serves to illustrate this point.

Consider the reasons why, as H&F point out, members of biological taxa have many highly correlated, and highly diagnostic, perceptual features. These properties of biological taxa result from the facts that (1) traits are passed reproductively from parents to offspring, and (2) living things are related to each other through descent from common ancestors. The more closely related two individuals are, the more traits they will share, and the more likely they are to share any given trait – a principle which holds across taxa for any arbitrarily selected trait, regardless of function (Sober 1988). For these reasons, phenotypic traits passed on through reproduction are very highly correlated in, and diagnostic of, biological taxa (especially lower-level taxa such as families, genera, and species). Not only does this mean that feature correlations can reliably be used to infer descent from a common ancestor, and thus shared taxonomic category membership, but they can also be used to induce the presence of a trait in one organism, given its presence in another of the same taxonomic category (e.g., if this mushroom is poisonous, another in the same genus is likely to be so as well) (Coley et al. 1997). From an evolutionary point of view, these are extremely useful and robust principles, and we expect evolved systems for representing living things to use them not merely by chance, but by design.

These principles do not hold, on the other hand, for artifacts. Artifacts are not the same kind of natural kind as are living things,

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

because they do not inherit their properties en masse through biological reproduction. Rather, they are endowed with their properties by the person who constructs them (Bloom 1996). We thus expect form-function correlations between exemplars of artifact categories (e.g., screwdrivers, chairs, wastebaskets), but only, as a matter of general principle, in those structural features directly related to function. For example, in order to function as a wastebasket, an object need only be of a certain size, with an opening, and capable of holding trash; it could be spherical or cylindrical; made of plastic, metal, or a hollow elephant's foot, and so on. From an evolutionary point of view, we expect artifacts to be represented much differently from living things, and for these representations to be constructed and accessed by different mechanisms. Indeed, the principles that such mechanisms rely on for making adaptive inferences may not hold across category types; for example, the assumption that traits will be richly correlated, regardless of function, holds for taxonomic living kinds but not for artifacts. This kind of functional incompatibility is what selects most strongly for the evolution of domain-specific mechanisms.

These points are not incompatible with the HIT model presented by H&F. Rather, they are meant to illustrate that one cannot merely stipulate that different information formats will be privileged for different categories of object. The model itself must specify the mechanisms or processes that produce such category-specific differences; if not, one is simply extracting from the model the assumptions that one puts in. Evolutionary theory, coupled with a careful analysis of the principles that have been reliably true of the kinds of objects with which humans interacted in the past, is an excellent source of hypotheses about the origins and nature of such mechanisms.

Limitations on current explanations of category-specific agnosia

Daniel Bub and Cindy Bukach

Psychology Department, University of Victoria, Victoria, B.C. V8W 3P5, Canada. **dbub@uvic.ca**

Abstract: The HIT framework accepts a number of assumptions that are widely held as plausible or even well established in the literature on category-specific agnosia. We point out that a number of these elementary conjectures, now almost taken for granted, have received little in the way of convincing empirical support.

Any interpretation of category-specific agnosia must attempt to define the nature of the difference between biological and man-made objects. Based on the evidence, there seems to be at least one class of patients that is affected by the degree to which members of a category share similarities in their appearance. These cases have difficulty identifying animals, fruits and vegetables, musical instruments, and so on, because within any one of these classes of objects, there are many different exemplars that overlap in both their perceptual and conceptual properties. Identification errors result when such a patient is presented with a particular exemplar (e.g., a lemon) due to failure in resolving the co-activation that occurs between similar members of a category (other fruits and vegetables; e.g., Arguin et al. 1996; Dixon et al. 1997).

This part of the story seems reasonable, and any integrative framework like HIT would have to include some way of producing decrements in performance when the system cannot fully disambiguate members within a class of similar exemplars (where similarity is defined along perceptual and conceptual dimensions). We may argue with the evidence from PET (Price et al. 1996) cited by the authors as support for their choice of re-entrant modulation as the mechanism by which a target object is differentiated from its competitors for perceptual identification (see Bub 2000 for a critique of the argument developed by Price et al.). But we agree with the general assumption that one form of category-specificity represents a failure to resolve competitive activation between the stored representations of similar objects.

How, though, can we explain the existence of patients who have more trouble identifying man-made than biological objects? We cannot appeal to notions of shared perceptual features or structural dimensions for most of these objects, so what is the underlying cause of the confusion that allows different birds or flowers to be identified, say, but not different kinds of tools or kitchen utensils? In our view, the idea that man-made objects require functional knowledge for their identification, and that selective impairment of such knowledge is responsible for certain kinds of agnosia, though intuitively appealing and almost universally endorsed in the literature, is actually accompanied by minimal evidence. Any account that simply reifies this assumption in a computational model, without evaluating alternative possibilities, can only provide an integrative framework of what we wish to conjecture but cannot prove.

The reasons behind our skepticism are as follows: First, while it is clear that we can easily state and intuitively understand that man-made objects have a particular use, it does not follow that this knowledge plays a direct role in establishing the identity of an object. Of course different objects are associated with different functions but do these functional differences directly influence perception? Written words evoke a strong and automatic internal experience of sound, yet empirical claims about the role of phonology in visual word recognition are based on much more than the fact that we hear words or subvocalize them when we read. The causal role of phonology in visual word recognition is difficult to establish unambiguously even after many years of careful research. Yet faced with the question of how different kinds of conceptual knowledge mediate word and object classification, we seem to have simply accepted, in the absence of adequate empirical support, the idea that because man-made objects conjure up ideas of their use, such knowledge must play a part in perception.

We regret that the functional imaging literature, on which Humphreys and Forde (H&F) place considerable emphasis, cannot be considered as a strong source of evidential constraint here. The problem, as outlined by one of us in a recent review (Bub 2000), is that the presence of activation suggesting, for example, that regions associated with hand actions are contacted during passive perception of tools does not permit an inference that such activation is needed for perceptual identification. By analogy, activation of articulatory regions is common during visual recognition, but we would not infer that this kind of representation is important for word comprehension. The activation of regions associated with hand movements could be a consequence of object recognition rather than a mediator of the process.

Finally, we note that the tacit assumption about functional knowledge in the neuropsychological literature is that it concerns information about the use of specific objects. But is impairment of such information the major cause of the failure to identify man-made objects? If so, we might expect to see confusion errors determined by functional similarity between exemplars; cup might be confused with mug, hammer with axe, screwdriver with wrench. Unfortunately, a detailed analysis of errors within the agnosia literature has yet to be undertaken, and indeed, the concept of functional relationship has been so variously defined that the question itself requires considerable clarification before an answer can be broached. It remains possible that some other type of relational constraint might better characterize the nature of identification errors for man-made objects.

Although linguists who specialize in semantic theory have generally devoted little attention to the representation of concrete objects, it is instructive to consider that Wierzbicka's (1985) analysis of man-made objects includes many contextually-relevant features, so that the situations in which objects co-occur is relevant to defining their actual purpose. The shape of a cup, for example, is not understandable without knowing something about how cups are used with saucers. And the difference in the appearance of a cup and mug is due to the fact that mugs are not used with saucers

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

and so do not taper at the base. Thus, man-made objects are related by contiguity rather than similarity of function. A saucer is not related to cup because they have similar functions nor is a knife similar in function to a fork, at least if we confine ourselves to questions of how each is used independently of the other. Their relationship exists because they are used together in many situations.

We conclude, then, that the issue of how man-made objects are related to each other so as to yield selective difficulty for certain kinds of patients remains a major question that has received little clarification as yet. The HIT framework accepts the conventional viewpoint that the function of an object interacts with its form. But the proof of this claim requires detailed predictions of the kind of objects that generate competitive activation during the process of identification. Is the competition between canoe and sailboat or between canoe and wigwam?

Beyond the sensory/functional dichotomy

George S. Cree and Ken McRae

Department of Psychology, University of Western Ontario, London N6A 5C2, Canada. **{gcree; mcrae}@julian.uwo.ca www.publish.uwo.ca/~gcree www.sscl.uwo.ca/psychology/cognitive/mcrae/mcrae.html**

Abstract: Most current theories of category-specific semantic deficits appeal to the role of sensory and functional knowledge types in explaining patients' impairments. We discuss why this binary classification is inadequate, point to a more detailed knowledge type taxonomy, and suggest how it may provide insight into the relationships between category-specific semantic deficits and impairments of specific aspects of knowledge.

Humphreys and Forde (H&F) propose a hierarchical model of semantic memory that distinguishes between perceptual structural and functional-associative semantic information. They use their model to explain patterns of category-specific semantic deficits, focusing on structure within perceptual representations, noting the top-down influence of semantic representations during processing. Unfortunately, limitations in existing theories prevent the authors from being able to provide insight, beyond the possible role of sensory versus functional knowledge, into the structure of the functional-associative component of their model. As they point out, definitions of sensory and functional knowledge have been underspecified, and tests of access to this knowledge in patients with category-specific deficits have been inconclusive. We agree that the notions of both sensory and functional features have been underspecified and inconsistent. We focus our discussion on possible improvements.

First, can a knowledge type account that goes beyond a binary sensory/functional distinction provide insight into reported patterns of deficits? To show that knowledge types can account for the basic behavioral phenomena, we (McRae & Cree, in press) used Wu and Barsalou's (in preparation) feature classification scheme (28 knowledge types) to classify each feature from our semantic feature production norms (549 concepts). We demonstrated that knowledge type construed in this manner accounts for several trends in the patient literature, including the tripartite distinction among creatures, nonliving things, and fruits and vegetables, as well as the fact that musical instruments tend to pattern with creatures, and nonliving food items tend to pattern with fruits and vegetables. Analyses showed that eight (rather than two) knowledge types explained the majority of the variance, suggesting that semantic memory is, indeed, organized by type of knowledge.

A critical prediction of the knowledge type account is that certain patterns of impairment (e.g., creatures being more impaired than nonliving things) imply accompanying patterns of impairment of specific feature types (e.g., impaired sensory features) that should differ by category in specific and predictable ways. Based on current theories that rely on binary distinctions, typical studies designed to test this prediction probe a patient's knowledge of perceptual and functional attributes of both living and nonliving things using feature verification, feature priming, naming from definition, or forced-choice from definition. To the best of our knowledge, none of these feature-based patient studies has adequately defined and tested the appropriate types of knowledge, nor adequately controlled the variables that would allow for valid inferences to be drawn from the results.

Such feature-based studies can be criticized for several reasons. First, a simple classification into sensory/functional knowledge ignores important distinctions in knowledge type, such as those demonstrated in our research. Studies often incorrectly treat all sensory knowledge as equivalent, comparing, for example, the parts of living things with the colours of nonliving things, as though these types of knowledge are a priori equally accessible (and central to the concept). Second, although it is well known that some knowledge types are more central/salient within a domain (e.g., the functions of nonliving things), tests often compare access to knowledge type across domains by distorting the definition of the feature type in the domain for which the knowledge is less central, and perhaps even absent (e.g., for creature functional features, researchers have mistakenly used items like squirrel <climbs trees>). Third, some types of knowledge are shared by many category members (e.g., creatures – <has eyes>), whereas the equivalent types of knowledge may be distinguishing in other categories (e.g., the external parts of nonliving things). Finally, it is essential to acknowledge the difficulty that exists in finding functional, or other "non-sensory" features of concepts, that do not have a significant perceptual component. A simple thought experiment in which one imagines how the feature in question was learned reveals this fact.

Given these constraints, how might one move beyond the sensory/functional dichotomy to test accessibility of types of knowledge in the three major domains? Table 1 demarcates the types of knowledge that our norms suggest may be suitable for such contrasts. Assuming that it is obvious why a comparison should work, we focus discussion on the comparisons that do not appear possible. First, fruits and vegetables have few external components, and share the ones that they do have, making it difficult to find similarly shared external components in the other domains. Second, external surface features, such as colour, tend to be consistent and informative for creatures and fruits and vegetables, but vary more widely for specific instances of nonliving things, making verification of these types of features easier (all else being equal) for the former. Third, internal components are rarely listed in norms for creatures, and even if reasonable estimates of the importance of this feature type could be derived, these features are highly shared when compared with the internal components of nonliving things. Finally, few creatures have explicit functions and associated actions, and those that do exist differ in kind from the functions and associated actions of fruits and vegetables and nonliving things.

It is important to note that although Table 1 offers potential comparisons, other variables exist that must be equated for valid inferences to be drawn. The feature verification literature with normals has shown that equating these variables is not easy, and certainly cannot be done based on intuition alone. In addition to those listed above, such as the number of concepts in which a feature occurs, some of these factors include: production frequency in a feature-listing task, ranked production frequency, cue validity, intercorrelational strength of a feature, number of features in a concept, concept typicality, and concept familiarity (Ashcraft 1978; McRae et al. 1999).

In conclusion, we agree with H&F that understanding category-specific semantic deficits requires integrating a number of factors. We add to their list of factors an account of types of knowledge that extends beyond the sensory/functional distinction. Providing evidence for or against such a knowledge type account will require significantly more careful research than has been conducted to date.

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

Table 1 (Cree & McRae). *Suggested comparisons for which tests of retention of specific types of knowledge may reveal insight into domain level deficits*

Knowledge type	Example features	C vs NLT	C vs F&V	F&V vs NLT	C vs F&V vs NLT
External components	Crab <has pinchers> Tricycle <has pedals>	✔			
External surface features	Frog <is green> Pear <is green>		✔		
Internal components	Cherry <has a pit> Car <has an engine>			✔	
Internal surface features	Corn <tastes sweet> Cake <tastes sweet>			✔	
Functions	Olive <used for oil> Bed <used for sleeping>			✔	
Actions	Strawberry <is picked> Spear <is thrown>			✔	
Participants	Pony <used by children> Doll <used by children>	✔	✔	✔	✔
Time	Pumpkin <used for Halloween> Cabin <used for vacations>	✔	✔	✔	✔
Affect/Emotion	Wasp <is annoying> Bomb <is frightening>	✔	✔	✔	✔
Contingency	Garlic <causes bad breath> Car <causes pollution>	✔	✔	✔	✔
Associated entities/ Objects	Lobster <eaten with butter> Saucer <used with tea cups>	✔	✔	✔	✔

Note: C = Creatures, NLT = Nonliving Things, F&V = Fruits and Vegetables.

Category-specific deficits: Will a simpler model do?

Jules Davidoff

Department of Psychology, Goldsmiths College, University of London, London SE14 6NW, United Kingdom. **j.davidoff@gold.ac.uk**
www.goldsmiths.ac.uk/academic/ps/welcome.html

Abstract: The purpose of the commentary is not to contradict HIT but rather to question whether its increase in predictive power outweighs the decrease in parsimony. For the refutable aspects of HIT, a simpler model for naming appears to achieve as much. Both models better fit the facts concerning naming performance than describe category-specificity.

When the number of reported cases of category-specific disorders was small, models to describe them could be accordingly simple. Heirarchical interactive theory (HIT) is, therefore, not an unexpected solution to the problem of providing some sort of unifying structure to the present complexity in the patterns of preserved and impaired deficits. The purpose of this commentary is not to contradict HIT but rather to question whether its increase in predictive power outweighs the decrease in parsimony. Our brain comes with multiple top-down connections. One does not doubt that they could serve as the neural basis for the re-entry processes required for HIT. However, at the present time a simpler model without feed-back loops contains the kernel of the explanation proposed in HIT.

The introduction of re-entry procedures and multiple knowledge bases in HIT promotes a danger "that the framework is in fact too powerful, and not open to refutation." In answer to the concern of irrefutability, Humphreys & Forde (H&F) point to two issues that make their model testable. The first would be if it turned out that "no fractionations occur between contrasting knowledge sources in different patients" and the second would be if it turned out that "naming could not be linked in a relatively direct way to the recruitment of visual knowledge over and above that found in categorization tasks." Of these, only the second is non-obvious.

As the refutable part of HIT is relatively circumscribed, it would repay considering whether a more simple scheme might not do just as well. Take a previous model (Davidoff & De Bleser 1993) (see Fig. 1) also based on the hierarchical structure of the logogen model (Morton 1979; 1985). The model in Figure 1 does not benefit from the Cascade procedures in Humphreys et al. (1988) and is clearly the worse for that. However, just as in HIT, the intermediate structures between the structural description and the phonological output lexicon are referred to as stores of knowledge. Though H&F do at various places go back to the older term of semantics, the change in nomenclature is to be recommended given the alternative uses of the term semantics in allied disciplines. More critical, H&F's model (Fig. 3), like that in Figure 1, proposes that it is only sensory knowledge that has a strong (direct) access to the phonological lexicon. However, it is no surprise that the knowledge required to carry out categorization tasks is not that used for naming. Access to associative or functional knowledge might be reasonably be expected to help recruit the names for the category (e.g., tool, furniture) or for where the objects are found (e.g., garage, living room) but it would be a non-obvious way to name the object. Thus, aphasic patients may show very good preservation of association knowledge combined with little useful naming (Laws et al. 1995).

Figure 1 collapses all forms of sensory knowledge. Clearly, dissociations within sensory knowledge do occur and the elaboration in H&F's Figure 3(Top) is warranted. However, the one form of

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

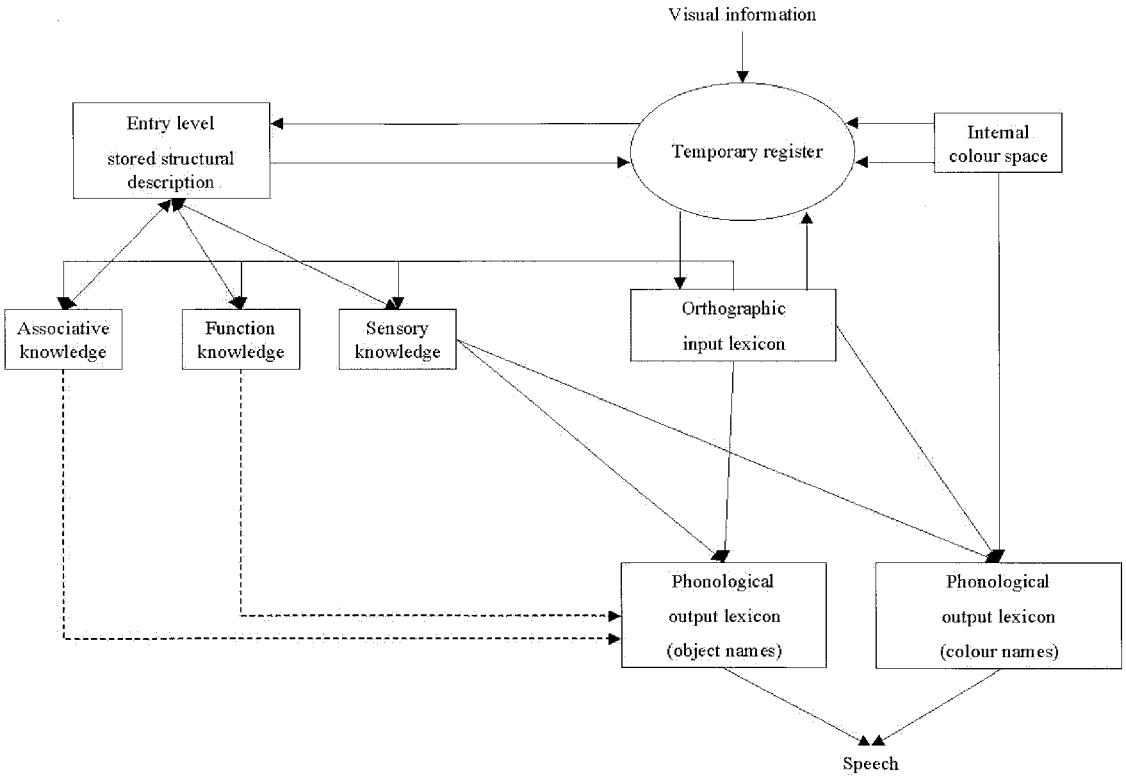


Figure 1 (Davidoff). A model of object naming (adapted from Davidoff & de Bleser 1993).

sensory knowledge that is not included in Figure 3(Top) is that of colour. Since Lewandowsky (1908) it has been known that an impairment in the retrieval of object-colour knowledge occurs without impairment for object identification. Thus, the position of this type of sensory knowledge is likely to be after structural descriptions and connected to the more direct path to naming (see Fig. 1). While the role of colour in superordinate categorization tasks is unsettled (Davidoff & Ostergaard 1988; Price & Humphreys 1989), its role in naming is undisputed (Biederman & Ju 1988; Ostergaard & Davidoff 1985; Price & Humphreys 1989). Moreover, the admittedly rather sparse clinical evidence is in favour of object-colour being beneficial to naming in aphasia (Bisiach 1966; Montanes et al. 1995). If colour is more reliably associated with living things, one might have thought that impairments in the retrieval of object-colour knowledge would have gone consistently with a category specific naming impairment. Yet, the evidence is that it does not (Forde et al. 1997; Luzzatti & Davidoff 1994). Indeed, more generally, the pattern of dissociations in patients shows no necessary link for functional or sensory properties to either living or nonliving things (Laws et al. 1995; Powell & Davidoff 1995).

H&F recognise that finding an explanation "concerning patients with visual/perceptual impairments that are not category-specific, is more difficult." Their answer is two-fold; H&F first argue for what is essentially a compensatory strategy that may be available to some patients by interrogating functional or associative knowledge. In Figure 1, it is clear such a procedure could promote an increase in naming ability. Second, H&F consider simulations with HIT using dynamic noise and note the consequent changes in naming output do not produce a category-specific impairment. However, there is a simpler solution. Sensory knowledge, even if reliably posted on the direct route to naming, is only of secondary importance. The critical aspect for identification, as H&F say, is sorting out the visual information at the level of the structural description. It is surprising how lacking in detail that selection process need be (Davidoff & Warrington 1999); their patient RK could accurately name animals but performed at chance in discriminating parts, global shape, and colours of the animals. Of course, there must be a limit to the changes made to an animal before recognition becomes impossible but with so much naming that can be achieved with so little, one may have difficulty in allocating any role for the other forms of sensory knowledge in explaining category-specific impairments.

The accumulation of data has not provided a clear answer as to why our minds so obviously divide the world into natural and artificial kinds. Figure 1 and HIT give a better account of naming than they do of category specificity. Natural kinds differ from artifacts in both their taxonomic classification (Disendruck & Gelman 1999) and in their uniformity of conceptualisation (Berlin 1999). Perhaps neuropsychology needs to look elsewhere for an answer.

Conceptual deficits without features: A view from atomism

Roberto G. de Almeida

Department of Psychology, Concordia University, Montreal, Quebec H4B 1R6, Canada. **almeida@alcor.concordia.ca**
www.psychology.concordia.ca

Abstract: Humphreys and Forde fail to account for the ontology of the "features" that they claim are constitutive of concepts. This failure is common to decompositional theories of conceptual representation. Category-specific deficits can be better explained by a theory that takes inferential relations among atomic concepts to be the key characteristic of conceptual representation and processing.

Concepts play a prominent role in the cognitive sciences because, qua mental representations, concepts are the very elements of thought and higher cognition. Regarding the nature of the represention of concepts, proposals range from several types of decompositional theories – from definitions to prototypes – to ver-

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

sions of atomism (Fodor 1998). In recent years, category-specific semantic deficits have been an important source of evidence for the nature of conceptual representation since patterns of deficits arising from focal brain damage or disease reveal aspects of conceptual organization which are not always transparent in studies with normal subjects.

But what, after all, is the nature of conceptual representation? If you start reading Humphreys and Forde's (H&F's) article with this question in mind – assuming research on conceptual deficits aims at unveiling the answer to that fundamental question – soon you realize that you are in a theoretical loop in the middle of an empirical maze from which you can leave only if you have answers to other even more fundamental questions: What is the nature of a perceptual feature? What is the nature of a functional feature? Are they concepts? Do they have the same status as other concepts (DOG, for instance; Is DOG a concept in the HIT model?)? Are they "sub-concepts"? Are they primitive elements? Where do they come from? Unfortunately, H&F do not have answers to these questions. However, the types of commitments one makes regarding the nature of conceptual representation set the conditions under which one can evaluate the pattern of conceptual deficits. In fairness, although H&F do not provide a full account of the nature of the elements constitutive of conceptual representation, it is quite clear that they are committed to the idea that concepts are represented by bundles of features.

Actually, it became a standard assumption in the field that no matter how one organizes one's conceptual stock – whether by hierarchical trees whose topmost branches are LIVING and NONLIVING or by any other taxonomic principle – concepts are represented by bundles of features. The idea is that the concept DOG, for instance, is in fact represented at some level as a finite or infinite set whose elements are things such as BARKING, FURRY, FOUR-LEGGED, PET, CANINE (see, e.g., Rapp & Caramazza 1991). The concept SCREWDRIVER, feature theories suggest, might also be represented at some level by a set of properties which should include things such as SCREW-DRIVER (function?), ELONGATED, HAS-A-HANDLE, and so on. Functional and perceptual properties – whether or not with different weights and whether or not represented within different subsystems – are in fact sets of features that supposedly contribute to concept tokening (i.e., for one's entertaining of the concept SCREWDRIVER in screwdriver contexts).

Elsewhere (see de Almeida 1999a), I have suggested that the cost of assuming that concepts are bundles of features is not only undermining the ontological foundations of one's theory but also committing to a noncompositional view of conceptual representation (see also Fodor 1998). In fact, in psycholinguistics, most empirical studies suggest that lexical concepts (i.e., concepts labeled by natural language morphemes) are not definitional or do not decompose into sets of more primitive elements (see, e.g., de Almeida 1999b; de Almeida & Fodor 1996; Fodor et al. 1975). This is certainly the case of verb concepts; and in fact, to my knowledge, thus far no one found evidence for "category-specific" verb concept deficits (e.g., that "features" such as CAUSE or GO, supposedly constitutive of complex verbs, are selectively impaired).

If not features, then what? In the remainder of this commentary, I will suggest that an atomistic view can better account for the pattern of dissociations of conceptual deficits. But before I move on to the analysis of category-specific deficits, I have to present briefly some of the basic assumptions of the atomistic-inferential view advocated here (for further discussion see de Almeida 1999a). First, it is assumed that concepts are atomic representations. Contrary to feature theories, the content of a concept is determined by epistemic liasons, not by its inferential relations (see also Fodor 1990). Thus, the content of X is determined by its property of "being an X," not by the relations obtained between constitutive elements (such as features). The present proposal assumes that all relations between concepts are inferential in nature – which means that for someone to possess concept X does not imply also that someone possesses concepts Y and Z: In this sense, concepts are individuated by virtue of their nomic relations with words, objects, events, and so on, not with their constitutive relations with other concepts. Second, it follows that any relations that are obtained between concepts X and Z are by virtue of their association or, more precisely, by virtue of their shared inferential domains – roughly, the sets of inferences unleashed by concepts. Much in the spirit of meaning postulates (henceforth, MPs; see Carnap 1959), those inferences are taken to be entailments. Suffice it to say for now that the inferential domain of concept X is a set of MPs. The assumption is that the inferential domain of X is the set A of inferences that are caused by X and also the subset B of inferences that are caused by Y but of whose entailments X takes part.

Categorical effects, thus, can be taken to constitute overlapping inferential domains. Under this view, category-specific deficits can be taken to arise from damage to concepts constitutive of certain MPs (for instance, the concept LIVING which might be related to ANIMAL via the MP [$\forall$x, ANIMAL(x) $\rightarrow$ LIVING(x)]). There are two related hypotheses to consider. The first is that the loci of the deficits are in the sets of inferences unleashed by the tokening of concept X (say, DOG). Since those inferences might be disrupted by the broken entailments that involve LIVING (as in the MP above), patients may have trouble selecting the appropriate lexical items in naming tasks. The assumption is that the tokening of DOG (by the word or the picture) causes the computation of the inferential domain of DOG – which intersects with the inferential domain of many other concepts, thus giving rise to semantic paraphasias. The second – and perhaps stronger – hypothesis places the locus of category-specific deficits in the selection of the appropriate concept, given a certain token stimulus. That is, it is possible that the specificity of the problem is in the causal link between the proximal stimulus dog and the concept DOG. Thus, in this sense, when dog is presented, DOG may be accessed but the inferences it unleashes (e.g., to ANIMAL) may lead the patient to consider other alternatives (CAT, COW) which are in the inferential domain of ANIMAL. Patients' strong performance in word/picture matching tasks even in "impaired" categories may provide evidence for the causal link hypothesis. Also, the fact that patients are able to produce items that are within the category (or, possibly, inferential domain) of the target item suggests that they are able to determine the nature of the stimuli (i.e., by assumption, dog may cause [DOG or CAT or COW]) but are unable to produce "dog" and in many trials they produce an incorrect response.

In sum, it appears that an atomistic-inferential theory of conceptual representation and processing can account for the pattern of dissociation without the perils of ontological vagueness and without sacrificing compositionality.

Structural descriptions in HIT – a problematic commitment

Markus Graf and Werner X. Schneider

Department of Experimental Psychology, Ludwig-Maximilians-University, Munich D-80802, Germany. **{grafm; wxs}@psy.uni-muenchen.de**
www.paed.uni-muenchen.de/mip/PSYCH/wxs/wwwdocs/index.html

Abstract: Humphreys and Forde conceptualize object representations as structural descriptions, without discussing the implications of structural description models. We argue that structural description models entail two major assumptions – a part-structure assumption and an invariance assumption. The invariance assumption is highly problematic because it contradicts a large body of findings which indicate that recognition performance depends on orientation and size. We will delineate relevant findings and outline an alternative conception.

We are in accordance with the two basic principles of the HIT model – a hierarchical processing structure and top-down (inter-

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

active) activation. However, one aspect of the HIT-model deserves closer consideration: Humphreys and Forde (H&F) employ the term "structural descriptions" to characterize the representations of the purely visual (i.e., "pre-semantic") system of object recognition. Thus, they use a technical term which refers to a specific type of recognition model, without offering convincing reasons for their choice and without discussing its implications.

Object recognition approaches can be classified into invariant property approaches, object decomposition approaches, and image-based approaches (see e.g., Ullman 1996; Wallis & Bülthoff 1999). In structural description models, objects are divided (parsed) into parts and represented in terms of their parts and the relations between the parts (e.g., Biederman 1987; Marr & Nishihara 1978; Sutherland 1968). Structural descriptions can be viewed as a mixture of part decomposition and invariant property approaches, whereby parts and their relations are defined as invariant properties (Ullman 1996). Therefore, we claim that structural description models include two major assumptions regarding object representations: (1) Object representations are characterized by a part-structure, and (2) are based on invariant properties and encoded in a propositional format, so that they are more or less invariant regarding spatial transformations (invariance assumption). Both assumptions are logically independent. We intend to show that the invariance assumption is problematic and doubt that Humphreys and Forde wanted to embrace it (see Humphreys et al. 1988, p. 68).

The authors barely explain their choice of structural descriptions, but instead refer to the cascade model (Humphreys et al. 1988) – which supplies no detailed explanation. Their justification of structural descriptions seems to be derived from what Humphreys et al. called structural similarity, a concept which is based on two measures: the number of shared parts (derived from attribute listings), and the ratio of outline contour overlap (see sect. 3.1). Outline contour overlap, however, is not specific for structural similarity, but is a general measure for visual similarity. Thus, only one of these measures relates to the notion of structural descriptions – and moreover, it only refers to the part-structure assumption.

Do object representations actually have a part-structure? There is empirical support for the hierarchical nature of object representations (Palmer 1977); and there is evidence that parts play a role in both real-time object recognition (Biederman & Cooper 1991) and in categorization (Goldstone 1996; Tversky & Hemenway 1984). However, these findings are not uncontroversial, as Srinivas (1993) could not replicate the results of Biederman and Cooper, using a similar procedure. Furthermore, other findings question a special role of parts in object representations (Cave & Kosslyn 1993; Murphy 1991). Overall, the empirical findings seem consistent with the first assumption of structural description models, but some doubts remain.

A crucial question for any recognition model is whether it can explain subjects' performance in recognizing objects after spatial transformations. A large number of behavioral studies accumulated which demonstrate that reaction times and error rates in object recognition depend systematically on stimulus orientation (both in the picture plane and in depth) and size (for reviews see Ashbridge & Perrett 1998; Jolicoeur & Humphrey 1998; Lawson 1999; Tarr & Bülthoff 1998). This pattern of results is not just due to early perceptual processes, but seems to be caused by high-level representations (e.g., Lawson & Humphreys 1998). Also neurophysiological data with single cells confirmed the orientation- and size-dependency (for review and discussion see Farah 2000). Moreover, an orientation effect was even found when all major parts or features of an object remained visible after a rotation in depth (Humphrey & Jolicoeur 1993; Lawson et al. 1994; but see Biederman & Gerhardstein 1993), or when elementary geometrical shapes – consisting of just one component – were rotated in depth (Tarr et al. 1998). Thus, effects of viewpoint are not just a result of self-occlusion of object parts or features. Taken together, the findings indicate that object representations are specific to orientation and size, which is regarded as evidence that object representations are image-based (e.g., Tarr & Bülthoff 1998).

This large body of findings cannot be explained with the existing structural description models, because they predict more or less invariant recognition performance – with minor exceptions like rotations in the picture plane (Hummel & Biederman 1992) and self-occlusion of parts in depth rotation (Biederman & Gerhardstein 1993). Is the inability to explain these data just a problem of the existing structural description models, or is it a matter of principle for this model type? The concept of a part-structure in visual object representations, by itself, does not necessarily lead to invariant recognition performance. The problematic point, however, is that structural description models are committed to the idea of invariance; they try to capture object invariance at the level of parts and relations. For example, size constancy can be accounted for by assuming that only relative sizes of the parts of a shape's contour are coded. It is this invariance assumption and the resulting abstraction regarding spatial transformations which leads structural description models to predict invariant recognition performance – and which makes them principally inappropriate to explain human object recognition.

How is it possible, then, to conceptualize a part-structure in object representations? It should be noted that evidence for hierarchical representations and the existence of a part-structure is compatible with most recognition models. Considering the orientation- and size-dependency of recognition performance, the most reasonable strategy is to integrate the concept of part-structure into an image-based framework of recognition. In such a model, relations between parts would be represented in an implicit way – and not explicitly as in structural description models. But this does not seem to be a disadvantage, because, to our knowledge, there is no empirical evidence indicating that relations between parts are coded explicitly in object representations.

Taken together, it seems advisable to refrain from making a commitment to structural descriptions, particularly because the invariance assumption of structural description models is highly problematic. We outlined an alternative image-based conception and suggest to use the more neutral term (structured) visual representations.

ACKNOWLEDGMENTS

This work was supported by a grant of the German Research Foundation, project "Task-dependent vision" (PR 118/19-1 TP7). We thank Adrian Schwaninger (University of Zurich) for helpful discussions.

Category-specific deficits and exemplar models

Koen Lamberts

Department of Psychology, University of Warwick, Coventry CV4 7AL, United Kingdom. **k.lamberts@warwick.ac.uk**
www.warwick.ac.uk/psych/lamberts

Abstract: Although Humphreys & Forde's HIT provides a comprehensive account of category-specific deficits, standard models of categorization and identification may also be able to explain many aspects of such deficits. The assumptions of an exemplar-based account of category-specific deficits are presented, and it is argued that exemplar models may be able to explain key findings on impaired object identification and categorization.

Humphreys & Forde's (H&F's) Hierarchical Interactive Theory provides a coherent and principled account of the bewildering array of findings on category-specific deficits. The model emerges from a research strategy that aims to unify a broad range of neuropsychological findings. This is an important strategy, and it is likely that HIT will become a benchmark for mod-

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

els of category-specific agnosia. However, it is somewhat surprising that well-established research on concept learning and categorization has had relatively little impact on theoretical and empirical work in neuropsychology. The basic mechanisms of categorization and identification are relatively well understood, but there have been very few attempts to apply standard models of categorization and identification to category-specific deficits. I would argue that this is unfortunate, for two reasons. First, standard theories of categorization can be used to derive precise, testable predictions about category-specific deficits, and thus form the basis of a productive research programme. Second, applying standard theories to category-specific deficits may reveal which aspects of processing might be altered in patients, without compromising the theory's ability to explain normal categorization or identification.

Recently, we have proposed that exemplar models of categorization and identification might form the basis of a comprehensive account of category-specific deficits (Lamberts & Shapiro, in press). Exemplar models, such as Nosofsky's (1986) Generalized Context Model, are among the best and most systematic accounts of categorization and identification across a wide range of stimuli and conditions. These models assume that category learning involves the storage of specific instances or exemplars in memory, and that subsequent categorization or identification is based on the similarity between the stimulus and the exemplars in memory. Exemplar models not only explain categorization and identification, but also the links between these tasks (Nosofsky 1987). Moreover, they form the basis of detailed accounts of the time course of category decisions (Lamberts 1995; 2000; Nosofsky & Palmeri 1997), and their principles underlie what is probably the most successful connectionist model of category learning (Kruschke 1992).

The exemplar-based account of category-specific deficits is based on the finding that living and nonliving objects tend to differ in perceptual similarity. H&F refer to several studies that suggest that living things are more similar to each other than nonliving things. This difference in similarity structure within the categories leads to very specific predictions about the nature of category-specific deficits. An exemplar account would predict that identification (which is the task of assigning a unique label to an object) is easiest for objects that have few similar neighbors. If discriminability of objects is low, or if exemplar memory is damaged (either of which could be the result of brain damage, see Lamberts & Shapiro, in press; Nosofsky & Zaki 1998), objects with many similar neighbors will be harder to identify than objects without similar neighbors. Applied to the living versus nonliving categories, nonspecific brain damage should lead to a selective impairment of identification of living objects, as observed in the vast majority of studies.

The exemplar account makes other predictions as well. If the task is categorization rather than identification (i.e., if objects have to be assigned to categories that contain more than just one member), nonselective damage should have the opposite effect: objects with similar neighbors in the same category should be categorized more accurately than objects without similar neighbors in the same category. This implies that brain damage should impair performance on nonliving objects more than on living objects if categorization is the task.

From this contrast, it follows that a critical test of the similarity-based account involves a direct comparison between identification and categorization of different objects. In one such study, Forde et al. (1997) carried out a number of experiments with their patient S.R.B., and found that his identification ability was impaired more for living objects than for nonliving objects. Forde et al. (1997) also examined S.R.B.'s ability to categorise living and nonliving things. He was shown line drawings of fruit, vegetables, animals, and tools and asked to classify them into their respective categories. S.R.B. scored very highly in this task. His overall pattern of performance was exactly as predicted by the exemplar account. Moss et al. (1998) also compared categorization and identification performance for their patient, R.C. Tested with the Snodgrass and Vanderwart picture set, R.C. was able to identify 50% of the pictures of artefacts, compared to only 9% of pictures of living things. In many cases in which R.C. failed to name the item, he was still able to provide some information about it. For 63% of the naming errors made on the living things in the test set, this included the correct superordinate name (e.g., animal for donkey, or fruit for peach), which indicates that his categorization abilities with these objects were relatively well preserved. Superordinate names were hardly ever produced for the nonliving things. Moss et al. (1998) also carried out a direct test of R.C.'s ability to categorise colour photographs of living and nonliving objects into their superordinate categories. R.C. was able to categorise the living things very accurately (93% correct), scoring within the normal range. However, his ability to categorise the nonliving objects (83% correct) was below the range for controls. Again, this confirms the exemplar model's predictions.

The exemplar account does have some characteristics that make it an attractive alternative for existing models of category-specific deficits. The model has been developed outside the neuropsychological literature, and has become one of the best-tested and most productive theories of perceptual categorization, identification and recognition. The model's simplicity and formal rigour are further assets. Of course, we cannot claim that exemplar models readily explain all aspects of category-specificity, but the models' scope and implications certainly merit further study.

Category-specific deficits: Insights from semantic dementia and Alzheimer's disease

Matthew A. Lambon Ralph[a] and Peter Garrard[b]

[a]*Department of Experimental Psychology, University of Bristol, Bristol BS8 1TN, United Kingdom;* [b]*The Institute of Cognitive Neuroscience, Alexandra House, London WC1N 3AR, United Kingdom.*

matt.lambon-ralph@bristol.ac.uk **p.garrard@ucl.ac.uk**
www.psychology.psy.bris.ac.uk.psybris **www.icn.ucl.ac.uk/**

Abstract: Recent investigations and theorising about category-specific deficits have begun to focus upon patients with progressive brain disease such as semantic dementia and Alzheimer's disease. In this commentary we briefly review what insights have been gained from studying patients of this type. We concentrate on four specific issues: the sensory/functional distinction, correlation between features, neuroanatomical considerations, and confounding factors.

Although many reports of patients with category-specific disorders have been based on stable brain-damaged patients (e.g., CVA: Sacchett & Humphreys 1992; HSVE: Warrington & Shallice 1984), it is only relatively recently that data from patients with progressive brain disease have been used to inform this issue. This is despite the fact that semantic impairments in progressive disorders such as Alzheimer's disease (AD) and semantic dementia (SD) were established many years ago (e.g., Warrington 1975). We summarise below some of the insights that have been gained from this source of data (for a fuller account, see Lambon Ralph et al. 1998b).

The sensory/functional distinction. Although there have been a number of reported associations between category-specific impairment for living things and relatively poor sensory knowledge (e.g., Gainotti & Silveri 1996), evidence for the complementary association (between nonliving things and functional knowledge) is sparse. One example was provided by a longitudinal analysis of definition and naming data of patients with AD (Lambon Ralph et al. 1997). Although the AD patients did not demonstrate category-specific differences in overall performance, Lambon Ralph et al. (1997) were able to show that a decline in the ability to name living things was associated with loss of sensory information whilst

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

poorer naming of artefacts paralleled degraded functional knowledge. As noted by Humphreys and Forde (H&F), there is an increasing number of single cases that call the causality of this association into question. There are patients with category-specific deficits for living things with equivalent sensory and functional knowledge. Patients with semantic dementia, in contrast, seem to show the opposite combination – although conceptual knowledge is degraded overall, their ability to give or confirm sensory attributes is particularly affected (Lambon Ralph et al. 1998b 1999). If one controls for concept familiarity, there is little evidence that this pattern leads to the predicted category-specific impairment (Bozeat et al. 2000; Garrard et al., in press a; Lambon Ralph et al. 1998a; 1998b; 1999).

With this issue in mind, we have recently completed analyses of a feature database collected from intact control subjects (Garrard et al. 2001). Despite the fact that Shallice and Warrington's (1984) original proposal was based upon assumptions about the nature of those features that distinguished individual concepts, neuropsychological assessments of feature knowledge have not differentiated between shared and distinctive attributes. In the feature database, we found no difference in the number of distinctive sensory features listed for living and nonliving concepts. The number of distinctive functional attributes was, as predicted, greater for artefacts. Further analyses also revealed another category difference that may prove to be critical. The greater visual overlap/similarity of living things noted by H&F, extends to conceptual representation themselves. We (see also McRae & Cree, in press) found that shared features were much more prominent for living than nonliving concepts.

Correlated features. Explanations of category-specific deficits that focus upon the importance of intercorrelated features predict that the direction of the category-difference should be related to severity (though the two best known positions make opposite predictions: Gonnerman et al. 1997; Tyler et al. 2000). Two recent studies including a relatively large number of AD patients found no positive evidence for either version of the theory (Garrard et al. 1998; in press c). The first, a cross-sectional analysis, found that the direction of the category-specific difference was not related to severity. Rather specific deficits for living or nonliving concepts were both more likely to be found in the most severely affected patients. More recently, after adding longitudinal data to the previous cross-sectional study, we found no positive evidence for the predicted cross-over in category-differences when plotted as a function of severity (either in terms of overall disease progression or the degree of semantic impairment). The lack of an effect in the target patient group (AD) may be explained by further analyses of the feature database, noted above. First, the number of significantly correlated feature pairs as a proportion of the total number of possible features pairs is extremely small, suggesting that any effect of feature co-occurrence is likely to be weak. Second, intercorrelation is confounded by feature distinctiveness – it is shared features that tend to correlate with others, again suggesting that it may be the distribution of shared versus distinctive features, which is critical.

Neuroanatomical considerations. Although it did not support the correlated feature accounts of category-specificity, the cross-section AD study (Garrard et al. 1998) did find positive evidence in favour of neuroanatomical influences. Those cases with relatively poor artefact knowledge/naming had more AD pathology in parietal regions, supporting the predominant view that there is a temporal versus frontoparietal difference at the heart of category-specific deficits. The semantic dementia cases are, again, a puzzle in this regard. The atrophy in these cases (e.g., Mummery et al. 2000) is focused upon the anterolateral aspects of the temporal lobes bilaterally and includes the inferior temporal gyrus. One might expect this to lead to relatively poor performance for living things as is found in some patients with HSVE for whom the distribution of pathology is somewhat similar (though it tends to involve medial as well as lateral temporal regions bilaterally: Gainotti et al. 1995).

Confounding factors. We finish where H&F began. Concept familiarity strongly influences the accuracy of patients with semantic impairment (Bozeat et al. 2000; Lambon Ralph et al. 1998a) and it certainly explains some, but not necessary all, of the underlying difference in category-specific cases (e.g., Funnell & De Mornay Davies 1996). We would argue that familiarity is no more a "nuisance" or confounding factor than different sensory-functional weightings of concepts – there is plenty of evidence that familiarity should be regarded as a critical part of understanding this issue. The magnitude of category-differences is reduced dramatically once familiarity controlled stimuli are used for assessment and there are at least three demonstrations of an interaction between category-specific differences and familiarity (Funnell & De Mornay Davies 1996; Gainotti & Silveri 1996; Lambon Ralph et al. 1998b). Implemented computational models of conceptual knowledge show that when familiarity is included during training, not only does it influence the models' overt performance (Lambon Ralph et al., in press), but also the nature of the derived semantic representations is changed. Familiar concepts tend to take up greater amounts of the semantic "space," making them much less vulnerable to simulated damage (Rogers & McClelland, submitted).

What is structural similarity and is it greater in living things?

Keith R. Laws

Department of Psychology, London Guildhall University, London E1 7NT, United Kingdom. **klaws@lgu.ac.uk**

Abstract: Humphreys and Forde (H&F) propose that greater within-category structural similarity makes living things more difficult to name. However, recent studies show that normal subjects find it easier to name living than nonliving things when these are matched across category for potential artefacts. Additionally, at the level of single pixels, visual overlap appears to be greater for nonliving things.

Two important and related issues in the paper by Humphreys and Forde (H&F) require examination: (1) What is the evidence that living things (i.e., with high structural similarity according to H&F) are more difficult to name than nonliving things (i.e., with low structural similarity)? and (2) What constitutes structural similarity? Does "contour overlap" adequately capture the degree of structural similarity within categories?

Turning to the first issue, H&F cite several studies to support the notion that normal subjects make more errors or are slower to name items with high structural overlap (exclusively living things: Gaffan & Heywood 1993; Humphreys et al. 1988; Lloyd-Jones & Humphreys 1997; Vitkovitch et al. 1993). All failed, however, to match across category for: visual complexity; some failed to also match for familiarity (Humphreys et al. 1988; Vitkovitch et al. 1993); or any variables including name frequency (Gaffan & Heywood 1993). Hence, the results of these studies may reflect the influence of the same artefacts that have proved troublesome in patient studies of category-specific effects. In contrast to the studies reported by H&F, several more recent studies have documented better (and faster) naming of living than nonliving things by normal subjects on sets of stimuli matched across category for familiarity, visual complexity, and name frequency. This has been demonstrated using a variety of paradigms including the rapid presentation paradigm of Gaffan and Heywood (Laws & Neve 1999) and the naming-to-deadline paradigm of Vitkovitch et al. (Laws 2000), naming latency (Laws 1999). These points raise two related issues that need to be addressed by H&F: first, that the results of unmatched studies are prone to artefact explanations and second, that when matching has been achieved, these studies contradict a central tenet of the

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

HIT (and the Cascade) model, that is, that living things (structurally similar according to H&F) are harder to name and identify than nonliving things (structurally dissimilar).

Turning to the more fundamental question: What constitutes structural similarity (at least for H&F)? In an attempt to quantify some aspects of structural similarity, H&F refer to two measures (from Humphreys et al. 1988), both of which may misrepresent the degree and/or characteristics of structural similarity. The first consists of subject ratings showing that living things have more shared parts than nonliving things. Nevertheless, common parts are often structurally quite dissimilar, for example, the ears of an elephant, rabbit, or dog, and so, may not provide a reliable indicator of structural similarity. The other, perhaps more pertinent, measure refers to contour overlap for items from the Snodgrass and Vanderwart (1980) corpus of line drawings. This measure may misrepresent "structural similarity" because it concentrates on gross contour overlap at the expense of internal detail. Direct evidence for this comes from Laws et al. (in press), who developed a measure of pixel-level overlap for subcategories of item from the Snodgrass and Vanderwart corpus. This measure was calculated by subtracting pixel values for all pairs of pictures (size standardized to 256-by-256 pixels, i.e., 65,536 pixels per picture) in any subcategory. The difference value was then squared and summed for all pixels in the array. Finally, the square root of the sum of squared differences gave the Euclidean Distance (ED) between the two pictures. Within-subcategory mean and standard deviation ED values were calculated for items in 6 living and 5 non-living subcategories (see Fig. 1).

By contrast with H&F, Laws and Gale found greater within-category structural similarity for nonliving than living things (ED was greater for living than nonliving things: $t = 3.9[df = 118]$ $p < .000$). In addition, they found that ED for body parts and musical instruments clustered with nonliving and living things respectively (see Fig. 1) and so, appears to incorporate the exceptions that occur in the category-specific deficits literature, that is, that musical instruments tend to be impaired along with living things and body parts with nonliving things. It seems that the measure of ED overlap has more justification than H&F's use of gross contour overlap (alone) since only the former incorporates both internal and contour information, predicts the performance of normal subjects with matched stimuli (outlined above), and may even account for the unusual performance of patients when naming body parts and musical instruments. To summarise, the points raised here seem difficult to accommodate within the HIT model proposed by H&F and indeed, point to some evidence that appears to oppose HIT predictions.

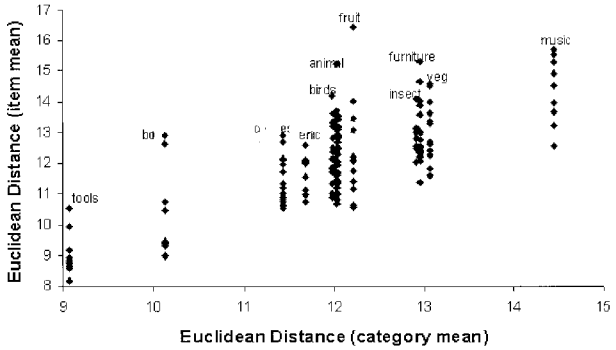


Figure 1 (Laws). Mean ED values for subcategories plotted against item mean ED. This shows that ED separates living and nonliving items and also clusters body parts along with nonliving things and musical instruments with living things.

What is specific about category specificity? Fractionating patterns of impairments and the spurious living/nonliving dichotomy

E. C. Leek[a] and E. M. Pothos[b]

[a]*School of Psychology, University of Wales, LL57 2AS Gwynedd, United Kingdom;* [b]*Department of Psychology, University of Edinburgh, EH8 9JZ Edinburgh, United Kingdom.* **e.c.leek@bangor.ac.uk**
www.psychology.bangor.ac.uk **e.pothos@ed.ac.uk**
www.psy.ed.ac.uk/Staff/academics.html#PothosEmmanuel

Abstract: What aspects of the data from studies of acquired category-specific impairments are relevant to theories of knowledge representation? Discussion in the target article focuses on the living/nonliving dichotomy. However, many case studies reveal considerably more complex patterns of impaired and preserved performance that undermine this distinction. We consider this evidence and discuss its implications for theories of knowledge representation.

The issue that we address concerns the kinds of data from studies of category-specific disorders that are typically deemed relevant to theories of knowledge representation. In their paper, Humphreys and Forde (H&F) focus their discussion primarily (but not exclusively) on what are described as cases of selective impairments affecting object knowledge for living versus nonliving things. Indeed, this apparent dichotomy is frequently presented in the literature on category-specific disorders as a key aspect of the data that theories of knowledge representation must explain. However, a closer inspection of evidence from single case studies suggests that category-specific deficits rarely break down in terms of this simple dichotomy. Rather, patients show complex patterns of impaired and preserved performance that vary among different categories of objects across the living/nonliving divide (Hillis & Caramazza 1991; Warrington & McCarthy 1987).

For example, JBR, one of the patients reported originally by Warrington and Shallice (1984), is described by H&F (and others) as having a selective deficit for living things. In fact, as H&F acknowledge, while JBR was poor at naming objects from biological categories such as animals and plants, he was also impaired at naming musical instruments, tools, fabrics and gemstones – a pattern seemingly at odds with the patient's supposed deficit for living things. YOT (Warrington & McCarthy 1987), when tested on an auditory/written word matching task, performed relatively well with objects from the categories of animals, occupations, vegetables, and fabrics, but poorly with flowers, furniture, and body parts. The patient PS (Hillis & Caramazza 1991) was disproportionately impaired at naming animals, birds, vegetables, and fruits, but had no difficulty naming varieties of foods and body parts. Another patient, JJ, described by the same authors, was good at naming animals and birds, but poor at naming across a diverse range of other categories including vegetables, fruits, body parts, clothing, transportation, and furniture. This pattern could be described as a case of the selective preservation of animals and birds, but not a selective preservation of living things, or a selective impairment for nonliving things. Other case reports describe patients with highly circumscribed deficits that seemingly affect a single category of objects or a small number of categories. These include selective impairments for body parts (Dennis 1976), proper names (Semenza & Zettin 1989), facial expressions (Rapcsak et al. 1989), foods (Warrington & Shallice 1984), indoor objects (Yamadori & Albert 1973), countries (McKenna & Warrington 1978), animals (Hart & Gordon 1992), as well as fruits and vegetables (Hart et al. 1985). Perhaps equally as striking are demonstrations of the apparent selective preservation of a single category of objects. For example, Shelton et al. (1998) have recently reported the case of a patient with a selective sparing of knowledge for body parts.

These observations seem to us to have important implications for theories of knowledge representation.

The spurious living/nonliving dichotomy. First, as others have pointed out, the complex patterns of performance that are ob-

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

served in patients undermine descriptions of the theoretically relevant data in this domain in terms of a simple living/nonliving dichotomy (Caramazza & Shelton 1998; Warrington & McCarthy 1987). Theories of knowledge representation that are constrained by patient data should not be restricted to (or even based upon) this distinction. Despite this, the dichotomy is still frequently used in the literature. Interestingly, as H&F note, fractionating patterns of category-specific impairments were one of the original motivations for the sensory/functional hypothesis (Warrington & McCarthy 1987) – with which Hierarchical Interactive Theory (HIT), outlined in the target article, shares several core assumptions. According to this view, categorical distinctions arise because of the differential weighting of modality-specific features in object representations (Warrington & McCarthy 1987). Although this proposal has been highly influential in accounts of category-specific deficits, attempts to formulate and test predictions of the sensory/functional hypothesis have largely focused on general claims about a putative asymmetry in the weighting of "visual" and "functional" features in the representation of living and nonliving things (Farah & McClelland 1991, but see Warrington & McCarthy 1987).

Fractionating patterns of performance and probabilistic models of knowledge representation. Second, if, as we argue above, discussions of category-specific disorders should not be based on a living/nonliving dichotomy, what is the appropriate range of facts from patient studies that should be used to constrain theories of knowledge representation? Should these theories be evaluated in terms of their ability to account for every statistically reliable pattern of impaired and preserved performance that is reported? Presumably, some aspects of the structure of knowledge representations (possibly including the distinction between biological forms and artefacts) are likely to be dependent on individual differences in experience, knowledge and expertise (Jackendoff 1990). Research on categorisation also shows that subjects notions of categories can be enormously flexible: Quite ad hoc groupings of objects can still make good categories if an appropriate context or theme can be provided (e.g., Barsalou 1985; Goodman 1972). Consequently, the types of conclusions that are drawn from patient data in this domain should also consider variation in the structure of knowledge representations across individuals. In approaches like HIT, there seems to be an implicit assumption that distinct patterns of performance across patients can be accounted for largely in terms of the locus and (presumably) extent of damage to the system. But it is also likely that impairments reflect individual differences, in addition to general constraints on representation. It seems to us that much of the interesting computational work in this domain must also take place at a more fine-grained level – that is, where the statistical parameters of intercorrelations among features underlying category structures are specified (e.g., Devlin et al. 1998; Leek, in press; McRae et al. 1997). But while H&F discuss possible intercorrelations between, for example, form and function in object representations (Moss et al. 1997), as currently articulated, it is not clear what HIT adds to this aspect of the debate.

The sensory/functional assumption or the data: Which do we keep?

Bradford Mahon and Alfonso Caramazza

Department of Psychology, Harvard University, Cambridge, MA 02138.
caram@wjh.harvard.edu **mahon@fas.harvard.edu**

Abstract: The HIT model explains the existence of semantic category-specific deficits by assuming that sensory knowledge is crucially important in processing living things, while functional knowledge is crucially important in processing nonliving things – the sensory/functional assumption. Here we argue that the sensory/functional assumption as implemented in HIT is neither theoretically nor empirically grounded and that, in any case, there is neuropsychological evidence which invalidates this assumption, thereby undermining the HIT model as a whole.

The HIT model of Humphreys and Forde (H&F) is a descendent of the sensory/functional theory (SFT) first proposed by Warrington, McCarthy, and Shallice (Warrington & McCarthy 1983; 1987; Warrington & Shallice 1984). Both the HIT model and the SFT explain category-specific deficits through the assumption that visual properties play a crucial role in recognizing and naming living things, while functional/associative properties play an equally important role in recognizing and naming nonliving things. We will call this assumption as instantiated in the HIT model, the sensory/functional assumption. Although this assumption has been challenged in the context of the SFT on both theoretical and empirical grounds (Caramazza & Shelton 1998) the authors of the HIT model worry that their theory may be too powerful, with too many dimensions, to be empirically refuted. In this commentary we will argue that this worry may be unnecessary: the evidence that has been used to undermine the SFT also undermines the HIT model.

One argument presented by Caramazza and Shelton (1998) against the SFT is that the putative distinction between the categories "living" and "nonliving" in terms of the relative importance of visual/perceptual and functional/associative properties is empirically untenable. When asked to list the important properties of a given exemplar, normal subjects do not list more visual properties for living things compared to nonliving things; nor do they list more functional/associative properties for nonliving things compared to living things (Caramazza & Shelton 1998; Garrard et al., in press; McRae & Cree, in press).

This objection may not apply as forcefully to the HIT model since it implements the sensory/functional assumption in a more complex way (even though it also assumes that visual properties are more important for the category "living things"). Within HIT, the principal aspect of the sensory/functional assumption is that members of the category "living things" are structurally much more similar to one another than are exemplars from the category "nonliving things." A consequence of this putative difference in perceptual "crowding" between categories is that any perturbation in processing visual structural representations would disproportionately affect the category "living things." A parallel argument is made for functional properties and the category "nonliving things." But what is the evidence that living things are structurally more similar to one another than are nonliving things?

Consider one result cited by H&F in support of the differential crowding assumption (and which presumably was used to estimate the structural similarity parameter in the modeling experiments they cite). Humprheys et al. (1995) found that exemplars from the category "living things" were rated to be more structurally similar to one another than exemplars from the category "nonliving things." The items used to obtain the ratings were 10 from the category "living things" (cat, tiger, cow, fox, horse, squirrel, donkey, camel, deer, and pig) and 10 from the category "nonliving things" (belt, blouse, boot, coat, dress, pants, ring, tie, watch, and sweater). However, the structural similarity ratings could easily go the other way, with nonliving things being rated as more structurally similar to one another than living things by the judicious selection of the items to be rated (e.g., living = {octopus, trout, seahorse, worm, dog}; nonliving = {screwdriver, awl, socket-wrench, file, chisel}). The issue of whether living things actually are more structurally similar to one another than are nonliving things remains an open question, which cannot be determined by rating only a few exemplars from the two categories, unless theoretically-unbiased sampling criteria are specified. But this has not been the case. Without theory internal motivations or independent empirical evidence, the authors' generalization from the twenty exemplars represented in the model to the entire categories of "living" and "nonliving" is dubious at best. In short, there is no compelling empirical evidence in support of the claim that visual properties play a greater role in processing the category "living things" than the category "nonliving things."

However, even if we were to grant for the sake of argument that living things are more structurally similar to one another than are

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

nonliving things, there are three types of neuropsychological evidence that are inconsistent with the sensory/functional assumption of the HIT model. These are the same empirical facts cited by Caramazza and Shelton (1998) to be inconsistent with predictions derived from the SFT.

(1) The categories of category-specific deficits are more fine-grained than would be expected if conceptual knowledge were organized merely on the basis of the visual/functional distinction. If we assume that visual properties are more important for the recognition of living things compared to nonliving things, we would not expect to find patients with a deficit to the category "animals" but not to "fruit and vegetables" or the reverse dissociation – impairment to the category "fruit and vegetables" with no impairment to "animals." But these patterns of category-specific deficit have been observed (Caramazza & Shelton 1998; Farah & Wallace 1992; Hart & Gordon 1992; Hart et al. 1985; Hillis & Caramazza 1991). The HIT model might account for this pattern of performance by making the ad hoc assumption that within the category "living," different types of visual properties are more important for some categories of exemplars (e.g., shape for animals) than others (e.g., color and texture for fruits and vegetables). However, it has already been shown that patients with a deficit in retrieving the color of objects do not show disproportionate difficulty for the category "fruits and vegetables" (Luzzatti & Davidoff 1994; Miceli et al., in press).

(2) Given the sensory/functional assumption, patients with a selective deficit to the category "living" should also display a greater deficit for the visual/perceptual properties of exemplars from the impaired category compared to their functional/associative properties. Caramazza (1998) reviews five patients (patient EW: Caramazza & Shelton 1998; patients GR and FM: Laiacona et al. 1993; patient JEN: Samson et al. 1998; Moss et al. 1998) with selective deficits to the category "living" but equal impairments for functional and visual knowledge of living things. H&F defend the HIT model against this pattern of deficit in two ways. First, they claim that there are two types of functional/associative knowledge. Functional knowledge of living things includes biological functions such as eating and the context in which the exemplar is found, while functional knowledge of nonliving things includes an object's usage and motor activity on the part of the user. They conclude: "Hence the 'functional' knowledge that can be impaired in patients with impairments for living things is not the same as the 'functional' knowledge about nonliving things that can be spared. Until . . . the type of information [in neuropsychological tests] has been matched across categories, there is no empirical reason to abandon an account which stresses the importance of different types of information for different categories" (sect. 4.3). It is not clear what kind of matching the authors are requesting. They could not be asking for task items to be matched for such factors as task difficulty or familiarity, because these variables were matched in the case of EW, for example. If they meant that the same items be included for the two categories, then by their own definitions of what counts as functional knowledge for living things and nonliving things, it is not possible to match the type of information between the two categories. It is not clear then what the force of their conclusion is supposed to be – it seems that they are building into their theory just the categorical distinction against which they are arguing. Are the authors claiming that functional knowledge is domain-specific?

Second, the authors propose that damage to visual/sensory features will simply make inaccessible the functional information associated with the visual representation. In support of this claim, they cite neuroimaging results in which an area assumed to be involved in storing information about the form of objects showed activation when subjects were answering functional questions about animals (Chao et al.,1999; Thompson-Schill et al. 1999; but see Caramazza 2000 for criticism of Thompson-Schill et al.). However, it is not possible to know whether a brain region that shows activation in an imaging study is necessary for the task in which the subject is engaged – the region could just as well be active because it is associated with the task. Because H&F are making a claim about the necessity of visual/sensory knowledge when answering functional questions about living things, the neuroimaging evidence is simply not capable of supporting their claim. Furthermore, the only implication that could be drawn from the neuroimaging results is that some visual knowledge is activated when subjects answer functional questions about animals. But the authors need much stronger evidence to support their quantitative claim that functional knowledge is impaired to the same degree as visual knowledge. Even if the neuroimaging data were capable of supporting a claim about the necessity of visual knowledge for answering functional questions, the data cannot be stretched into supporting a claim of why patients show equivalent deficits to visual and functional knowledge.

Consider then the state of the HIT model's explanation of the performance of patients EW, GR, and FM (three of the patients cited above who had deficits to the category living and equal impairments to visual and functional knowledge for that category). These three patients were at close to normal levels of performance for visual and functional knowledge for the category "nonliving." HIT must at the same time make the claim that damage to visual knowledge is extensive enough so as to create equivalent impairments for functional knowledge in the category "living," and acknowledge that visual knowledge is unimpaired for the category "nonliving" in the same patients. Thus, we are left with the unanswered question: How is it that visual/perceptual knowledge of nonliving things is unimpaired in these patients? Are we to entertain the claim that there are two kinds of visual/perceptual knowledge, one kind for living things and one kind for nonliving things?

Neither can the performance of these patients be dismissed by the claim that ceiling effects in the unimpaired category are masking a deficit. For example, patient GR (Laiacona et al. 1993) was impaired for both visual and functional knowledge of living things (55 and 58% correct, respectively) and yet was not at ceiling for visual and functional knowledge of nonliving things (91 and 84%, respectively). Thus, it cannot be the case that a deficit for the visual properties of nonliving things is not being detected.

(3) Even if we were to grant the claim that visual knowledge is necessary in order to answer functional questions about animals, the central prediction made by this hypothesis is not borne out. The prediction is the same as that made by the sensory/functional assumption: patients who have damage to visual/perceptual knowledge must have a deficit to the category "living." However, patients have been reported with an impairment to visual/perceptual knowledge but no deficit to the category "living" (Coltheart et al. 1998; Lambon Ralph et al. 1998). H&F acknowledge the problem posed by this last pattern of deficit and present an empirical argument in defense of the HIT model in which they cite an unpublished simulation (Watson & Humphreys, in preparation, cited in Humphreys & Forde 2000) which supposedly shows that it is possible to damage visual knowledge in the model in such a way that the "differential impairment to living things" is decreased. However, evidence from simulations in which the difference in impairment levels between categories is reduced when certain types of "lesions" are made to the model is not sufficient to save HIT. What they must show is that the model performs in a quantitatively comparable way to the behavior of patients: severely impaired knowledge of visual/perceptual features but worse performance in naming exemplars from the category "nonliving" than "living" (Lambon Ralph et al. 1998).

It does not seem that theories based on the sensory/functional assumption are able to account for the facts of category-specific deficits. Thus, in response to the authors' worry that the HIT model's current articulation is too powerful to be empirically refuted, we have argued that there already exists sufficient neuropsychological evidence to invalidate the model's core assumption, and consequently, the model as a whole.

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

On disentangling and weighting kinds of semantic knowledge

Agnesa Pillon and Dana Samson

Unité de Neuropsychologie Cognitive, Université Catholique de Louvain, Louvain-la-Neuve, 1348, Belgium. **{pillon; samson}@neco.ucl.ac.be**
www.neco.ucl.ac.be

Abstract: To account for category-specific semantic deficits, Humphreys and Forde propose to fractionate semantic memory into multiple sensory and functional knowledge stores. There are reasons to doubt the empirical productivity of this proposal, unless theoretically motivated principles of distinguishing and weighting the different kinds of object knowledge can be spelled out in detail.

In order to account for category-specific semantic deficits, Humphreys and Forde (H&F) assume that different forms of sensory and functional knowledge are weighted differently in the representation of different categories of objects and that each of these knowledge forms is represented within a separated knowledge store. Our commentary focuses on this particular aspect of the Hierarchical Interactive Theory. We think further specifications are needed before this Multiple-Knowledge-Stores (MKS) account for category-specific semantic deficits could be put into the appropriate empirical test.

First, the account should make more explicit the principles that are assumed to organize knowledge in semantic memory. Unless these principles are spelled out in detail, there is no theoretically-motivated basis on which a consistent taxonomy of the different kinds of object knowledge could be drawn and the relevant tests for assessing patients' semantic knowledge be designed. There is no obvious means to relate a given property of an object to one or another kind of knowledge. How could we decide, for example, that "kangaroos move by jumping" is a property pertaining to knowledge of "biological functions," "action," or "motion" and on which basis could we assume that knowledge of this property relies on a functional/biological, functional/action or visual/motion knowledge store? The present ambiguity comes in part from the knowledge stores being equated with "perceptual recognition" stores. On one hand, this invites the interpretation that knowledge of object properties is assumed to be organized according to the sensory modality or modalities that are used to learn and transact particular objects (Tranel et al. 1997). On this interpretation, knowledge of biological functions, as well as knowledge of characteristic motion (for animate and inanimate objects) and, even, of object usage, might be considered as being heavily dependent on *visual* knowledge. On the other hand, a number of non-incidental aspects of object knowledge seemingly could not be linked uniquely or directly to a specific perceptual or motor knowledge store. We know where and how animals live, how to eat and cook vegetables, how we have to take the bus, and that we can use a kettle to boil water for tea. Such knowledge seems to require the various perceptual and motor features on which objects and actions are based being integrated within a relation specifying the specific space, event, process, or goal in which they are involved. It is unlikely that the elementary sensory and motor features that make up the vocabulary of perceptual and motor systems could suffice to represent such knowledge. Integration of features might, at the very least, require some inter-modal or supra-modal processing system. It is unclear which knowledge store, located at which hierarchical level of the proposed architecture, could achieve this goal.

Second, the account should clarify what is meant by the different forms of knowledge being *weighted* differently in the representation of different categories of objects. Weighting could be conceived of as a quantitative or a qualitative feature. In the former case, it could refer to the ratio of the number of a given kind over other kinds of properties (Farah & McClelland 1991) or to the relative prevalence of the different sensory modalities in transacting objects (Tranel et al. 1997). In the latter, one could consider the extent to which various kinds of properties are important in discriminating between two similar entities (Warrington & McCarthy 1983), the relative contribution of each kind of property in categorizing an exemplar as being or not being a given object (Malt & Johnson 1992) or the relative centrality (i.e., causal status) of different types of properties within the relational structure of the object's features (Ahn 1998). Adopting one or another view of weighting might have significant consequences. Empirical estimates based on different views did produce contrasting results. For example, Farah and McClelland (1991) found that visual properties have greater weight than functional properties for living but not for nonliving things, whereas Ahn (1998) found that both visual and functional properties could be causally central in the representation of living and nonliving things.

Hence, in its present state of specification, the MKS account does not allow to empirically derive a principled estimate of knowledge weighting across categories, which is crucially needed yet to draw new predictions about which subsets of living or/and nonliving objects could selectively be impaired, given hypothetical damage to one or another knowledge store.

Still, even such estimate were available, the cascaded and interactive processing framework of the MKS account makes it difficult to predict, without the help of a simulation, which pattern a patient with a given category-specific semantic deficit should display when answering questions about different kinds of object properties. Let us suppose, for example, that knowledge of fruit is highly dependent on color knowledge while knowledge of animals is highly dependent on shape. One could not rule out the following expectations: selective damage to the color knowledge store should impair naming of fruit and, at the same time, the retrieval of color, shape, and function of fruit, because color knowledge might also be required when answering questions about shape and function of fruit – while, for animals, naming as well as retrieving knowledge of shape, function, *and* color could be spared because, in this case, access to color knowledge might be supported by the integrity of activation from spared knowledge stores (shape and function). Hence, the question is raised whether a localised damage to such an architecture could eventually result in the pattern that is intuitively expected given selective damage to a knowledge store, that is, a selective semantic deficit for a category (e.g., fruit), associated with a *selective* impairment in accessing one kind of knowledge (color) across all categories (fruit, animals, etc.). If this intuitive prediction were shown to be wrong, then we ask the question of how evidence could be sought in a patient's performance for the hypothesis that its category-specific deficit originates from a specific knowledge store being *selectively* damaged.

ACKNOWLEDGMENTS

We are grateful to Bruno Rossion for helpful comments on an earlier version of the text. Agnesa Pillon is Research Associate of the National Fund for Scientific Research (Belgium).

Mutual access and mutual dependence of conceptual components

Friedemann Pulvermüller

MRC Cognition and Brain Sciences Unit, Cambridge CB2 2EF, United Kingdom. **friedemann.pulvermuller@mrc-cbu.cam.ac.uk**

Abstract: The HIT model comes close to a view suggested by Donald Hebb, that cognitive representations are organized as distributed neuron webs, cell assemblies, whose components are mutually connected and whose internal connections provide continuous information exchange among sub-components of the representation. Two questions are asked related to (1) the organization of internal connections of a concept representation and (2) the conditions under which information exchange between components are assumed in the HIT model.

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

Humphreys and Forde's (H&F's) seminal paper summarizes a wealth of neuropsychological and neuroimaging data on category-specific brain processes. Their synopsis is convincing. Concepts are laid down in the human brain as distributed representational and access systems. The cortical topologies of the network organizing an individual entry in semantic memory reflects linked sensory knowledge (visual, acoustic, gustatory, etc.), action knowledge, and verbal attribute. Lesions in specific parts of the distributed networks can affect conceptual categories to different degrees (e.g., action-related concepts more than visually or related ones), and at different processing levels (e.g., in object perception, processing of functional attributes, or naming). The motor and sensory modalities crucial for defining a concept, together with the degree of feature overlap of the concepts, are crucial for determining category-specific brain processes.

This approach exhibits family resemblance to other recent neuroscientific models of language and cognition which agree on two points.

1. A word and the concept it stands for are laid down in the brain by a distributed neuronal representation. The distributed representation consists of sub-parts, functionally distinct *components,* representing the concept's perceptual attributed (e.g., structural descriptions of visual attributes), its action attributes (sometimes referred to as "functional" attributes), and its word form (*functional dependence postulate*).

2. There are reciprocal connections allowing for reciprocal activity flow between the sub-parts of a representation. The bidirectional connections can yield interactive processing of information of different types (e.g., phonological and semantic) and mutual functional dependence (*functional dependence postulate*).

The new neuroscientific models of language and cognition are distinct from classical modular proposals denying functional dependence (2). A modular view may postulate that information at each level (perception/structural, action/functional, word form) is autonomous (or "encapsulated"). They are also different from accounts questioning functionally distinct components (1), for example, models implying that all conceptual and meaning-related items are stored in one uniform conceptual layer.

Although the new models exhibit a family resemblance to each other, it may be fruitful to look at possible differences. Le me ask two questions which occurred to me when comparing H&F's theory to a model we proposed (Braitenberg & Pulvermüller 1992; Pulvermüller 1991; 1992; 1996; 1999) which builds upon Hebb's (1949) concept of cell assemblies.

Paths through the networks. Along which routes would information travel through the neuron networks representing concepts and words? H&F present an architecture where structural descriptions are calculated from visual input. These are mapped onto functional knowledge including action knowledge and onto other knowledge types. Only from there is the word form knowledge in the phonological output lexicon accessed (target article, Fig. 3A; see also Fig. 1A of this response). A dog picture reminds one of being chased by such an animal (and of its various other attributes) and finally calls up the word "dog." If a concept is accessed through a different input modality, for example, if a guitar is recognized from its sound, an analogous process activating the word form representation indirectly, after excitation of structured descriptions of visual features or action associations (cf. Fig. 3B). The guitar sound would trigger the mental image of the instrument and possibly the finger movements necessary for playing it, and this indirectly triggers the mental pop-up of the word.

As an alternative, the neuron populations organizing the word form would be connected directly to the neuronal equivalent of both the structural and functional knowledge (Fig. 1B of this commentary). Direct links would exist from perisylvian language areas, where word forms may be housed, to the visual areas in inferior temporal lobes, where the structural visually-related information may reside, and to the action-related areas in the frontal lobes, where certain types of functional knowledge may be expected. The visual picture of the dog would activate the word first, and finally activity would spread to other areas where, for example, action associations or other perceptual attributes are stored. In an analogous manner, the sound of the guitar would first awaken the word which, in turn, would activate the shape.

As a third possibility, it may be that all relevant components of a neurally stored concept are directly linked by reciprocal long-distance connections. For example, the structural visual representation in inferior-temporal and/or occipital areas would be linked to the word form representation in the perisylvian language areas which, in turn, would connect to the functional representation in action-related fronto-central areas, which, finally, would project to, and receive projections from, the visual areas (Fig. 2 of this commentary). This architecture may be considered a synthesis of the earlier proposals and would be consistent with neuroanatomical knowledge and with the observation that differential activation seen close to action-related cortical areas occurs quite early, within the first 1/4 second after presentation onset of written words, around the time the lexical access process may take place (Pulvermüller et al. 1999; 2000).

Mutual dependence of components. Of utmost importance is H&F's argument that re-entrant perceptual information from the visual modality may be essential in object naming (target article,

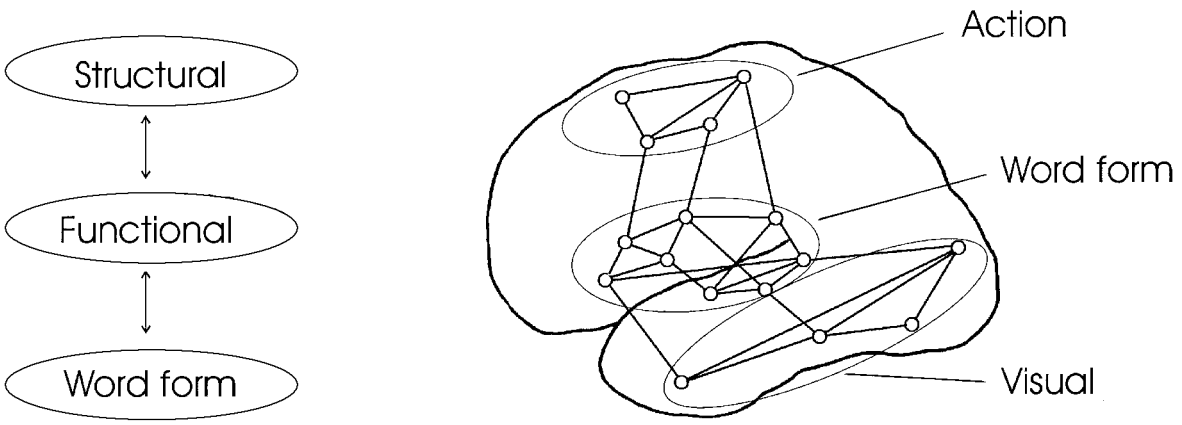


Figure 1 (Pulvermüller). A: Sketch of the functionally effective connections illustrated in Humphreys and Forde's HIT model. Information transfer is from structural visual representations to other sensory and action-related functional representations and finally to word form representations. B: An alternative view may be that there is direct access of word form information from all other functional components, so that, for example, both action-related and visually-related information can directly interact with the word form represent action.

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

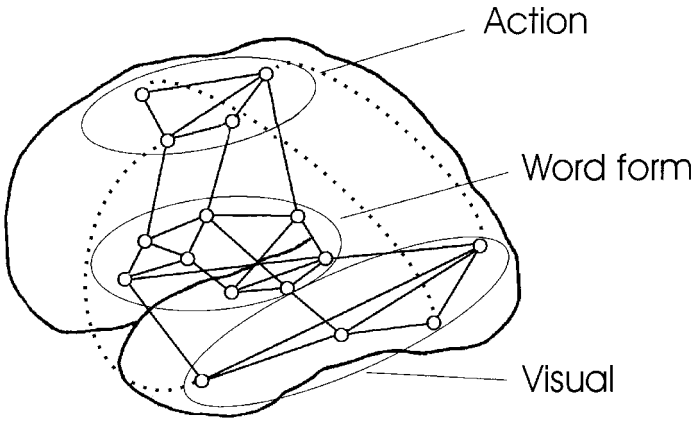


Figure 2 (Pulvermüller). The proposed synthesis is that all sub-components of a distributed cortical representation of a concept are related, for example, to its visual and action attributes and to its word form.

sect. 6). We recently reported a complementary finding, that action-related brain processes may be relevant for the access to words that refer to actions. Patients with lesions in their right hemisphere not dominant for language were investigated. All suffered from left-sided hemiparesis caused by lesions destroying part of the motor cortex and adjacent prefrontal areas of the right hemisphere. These patients participated in a speeded forced-choice lexical decision experiment where they had to indicate by button presses whether a briefly presented letter string was a correct German word or a pseudoword they had never encountered. Although the patients did not show language difficulties in everyday life or on clinical aphasia tests, they exhibited category-specific processing deficits for action verbs, that is, reduced accuracy (and prolonged latencies) of lexical decisions on action verbs relative to matched word stimuli not primarily characterized by action associations. We interpret this finding as evidence for the importance of re-entrant action-related information to the word form representation in the lexical decision task. Taken together, these data are consistent with the idea of functional inter-dependence of the sub-components of conceptual-linguistic representations (cf. point (2) above) (Genkinger & Pulvermüller, in press; Pulvermüller et al. 1998).

H&F postulate "that different forms of knowledge are activated in a reiterative fashion for name retrieval to operate" (sect. 6.2, para. 7). It is, however, not clear to me how this generalizes to other cognitive operations and tasks. A more radical view may be that such continuous re-entrant information flow between processing components and cortical areas is a fundamental characteristic of cortical processing. Would the authors consider this a possibility?

Further evidence in support of a distributed semantic memory system

Eleanor M. Saffran[a] and H. Branch Coslett[b]

[a]*Department of Neurology, Center for Cognitive Neuroscience, Temple University School of Medicine, Philadelphia, PA 19140;* [b]*Department of Neurology, University of Pennsylvania School of Medicine, Philadelphia, PA 19104.* **saffran@astro.temple.edu** **hbc@mail.med.upenn.edu**

Abstract: We offer additional points that support a distributed semantic memory: (1) the activation of representations that are modality-specific; (2) patients with inferotemporal lesions fail to activate visual object representations in semantic tasks, although normal subjects do; (3) direct activation of action systems from pictorial information, but not from words; (4) patients who demonstrate superiority with abstract words fail to access perceptual representations.

We endorse Humphrey and Forde's (H&F's) proposal for a distributed semantic memory system, a position we have advocated elsewhere (Breedin et al. 1994; Saffran 2000; Saffran & Schwartz 1994; Saffran & Sholl 1999). In this commentary, we would like to raise several additional points, not cited by these authors, that provide further support for this view:

Dependence on an abstract, modality-free mode of semantic representation, as advocated by Caramazza et al. (1990) and others (Lambon Ralph & Howard 2000), is not sufficient. One of the strong arguments in favor of this position is that it is parsimonious. But it may be questioned whether the brain respects the dictates of parsimony. Note that in the case of the visual system, different tasks are apportioned to different areas of the brain (e.g., object recognition, the representation of color information, motion, spatial location, etc.). Furthermore, evidence from neuropsychology strongly implicates the need to activate representations in specific brain areas. Consider, for example, the impairment known as "autotopagnosia." In these cases, which involve damage to the left parietal lobe, patients are unable to point to parts of the body, even though the words themselves may sound familiar (e.g., Buxbaum & Coslett 2000). In contrast, these individuals are able to point to parts of objects, or even parts of non-human (e.g., horse) bodies; additionally, these individuals may be able to point to the body part on which an item of jewelry is worn (e.g., watch), despite the fact that they are unable to point to the same body part on command. These data suggest that lexical-semantic knowledge of the human body is distinct from representations underlying knowledge of objects. Consider color information: could a patient have knowledge of the color RED (even if she knew that cherries are red) if there were no pointer to information regarding that color in the visual system? The evidence suggests that she would not (e.g., Riddoch et al. 1988). And there is the recent finding that a patient with bilateral damage to the amygdala, who was capable of recognizing photographs of people, was able to recognize expressions of happiness, sadness, and disgust. But this same individual had difficulty recognizing angry facial expressions, and most difficulty describing fearful facial expressions as afraid (Adolphs et al. 1994). These observations are consistent with a model that requires the activation of information represented in systems in the brain that respond to and store specific types of information.

There is evidence from a functional imaging study that normal subjects activate posterior temporal regions concerned with object recognition in semantic tasks, whereas semantic dementia patients do not (Mummery et al. 1999). These patients have incurred damage to the anterior, inferior temporal lobes, the same area that is affected in herpes simplex encephalitis patients, who constitute the majority of patients impaired on living things. These findings suggest that such patients are truly impaired in accessing percep-

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

tual information about objects. Note that the information may not necessarily be lost, but merely inaccessible to other information stored in semantic memory, as in a case (DM) we have reported (Breedin et al. 1994; Srinivas et al. 1997). As noted by H&F, DM performed rather well on object decision tasks, and demonstrated priming for perceptually similar objects, although he provided little or no information when shown pictures of objects (for example, cued by the letter /b/, he called a picture of a bone "a buffalo"). But this impairment did not prevent him from carrying out complex, sequential tasks utilizing objects, such as packing a lunch box or wrapping a present (Buxbaum et al. 1997). A similar patient, reported by Lauro Grotto et al. (1997), was able to prepare a variety of foods, even though she could not associate the labels for those foods with the relevant objects. The performance of these tasks may be mediated by connections between object recognition areas in the inferotemporal lobe and action-related knowledge stored in the parietal lobe.

As just noted, there are cases in which patients are capable of carrying out complex tasks with objects, even if they perform poorly on tasks that assess their knowledge using pictures of objects (e.g., the Pyramids and Palm Trees test of Howard & Patterson 1992). This suggests that there may be a privileged relationship between visual representations of objects and action, as earlier suggested by Gibson (1979). Data from a study by Rumiati and Humphreys (1998) support this position. In studies with normal subjects, they found that action errors were likely to occur with visual stimuli (e.g., mistaking a hammer for a razor and executing a shaving motion), whereas semantic errors were more likely for verbal stimuli. In our laboratory, we have found that pictures of manipulable objects are more likely to elicit verbs as associates than words do (Saffran et al. 2000). Recent studies with optic aphasics, who cannot provide names for visual representations of objects (although they can name to tactile presentation and verbal description), have shown that these patients can provide names for visually presented actions (i.e., verbs) (Campbell & Manning 1996; Teixeira Ferreira et al. 1997).

Finally, there are a number of reported cases of patients who demonstrate loss of the normal concreteness effect, namely, better performance with concrete than abstract words (e.g., Paivio 1991). There is evidence that these individuals, who do better on abstract words, fail to gain access to perceptual information about objects, knowledge that is critical to the meanings of concrete words (Breedin et al. 1994; Marshall et al. 1996). Moreover, five of the seven such cases reported in the literature (two were not tested on this contrast) performed better on artifacts than living things, another finding that supports the perceptual deficit hypothesis as the source of the impairment with biological objects (Saffran & Sholl 1999).

These findings argue persuasively in favor of a distributed system for semantic representation. It appears that the brain is not a parsimonious instrument, but rather one that assigns specific operations to particular subsystems. While it is likely that there is, in addition, a modality-free system for storing information about the world, it appears that this knowledge base is not sufficient to support normal performance on semantic memory tasks.

Where are object properties? In the world or in the mind?

James Tanaka

Department of Psychology, Oberlin College, Oberlin, OH 44074.
tanaka@cs.oberlin.edu **www.oberlin.edu/psych/labs/viscoglab.html**

Abstract: This commentary questions whether the category properties of an object can be determined independent of the experience of the categorizer. Expertise studies have shown that the judged properties of an object can differ from expert to novice and from expert to expert. The expertise findings indicate that object properties exist not only in the world, but in the mind of the categorizer.

Among the most compelling puzzles in neuropsychology is how brain damage can selectively impair knowledge about one category (e.g., living things) while leaving knowledge about another category (e.g., nonliving things) relatively intact. What do these category-specific breakdowns tell us about the neurological organization of object knowledge? Theories, such as Humphreys and Forde's (H&F's) HIT model, propose that the semantic system partitions object knowledge according to its sensory (visual, tactile, auditory) and functional properties. To account for category-specific impairments, such as breakdowns between living and nonliving things, property-based models propose that objects from the category of living things contain a disproportionate amount of visual features in contrast to nonliving objects which possess a greater number of functional features. Consequently, damage to visual properties will more severely impair one's knowledge of living things, whereas damage to the functional properties will more greatly impair the knowledge of nonliving things. While property-based models provide a viable account of category-specific impairments, these models leave unspecified exactly how the object properties of a category – sensory or functional – are derived. These models assume that object properties exist somewhere "out there" in the world, waiting for the human categorizer to discover them. In this commentary, a contrasting viewpoint is presented where the pragmatics of object categorization are stressed. According to the pragmatic view (see also Schyns 1998), object properties are not only determined by the intrinsic features of an object (e.g., color, shape), but also by the experience and motivations of the categorizer.

How might experience influence the specification of object properties? Methodologically, this question can be addressed by asking experts to list the properties of objects in their domain of expertise and compare these features to the properties listed by novices. When Marjorie Taylor and I asked expert dog judges and bird watchers to list properties for dogs and birds (Tanaka & Taylor 1991), we found that the expert features differed in number, proportion, and type from the features listed by novices. Overall, experts mentioned a greater number of functional features for birds and dogs than novices. As pointed out by H&F, "functional" with respect to living things typically refers to the behavioral properties of the object (as opposed to instrumental properties related to the function of nonliving things). For example, bird experts knew more than novices about migratory patterns and feeding habitats of birds, whereas dog experts were more knowledgeable about the physical abilities and temperaments that distinguish particular breeds of dogs. These findings suggest that the functional properties of an object are not invariant, but emerge as the result of the everyday interactions between the object and the categorizer.

If expertise influences a category's functional content, it is possible that expertise might also influence the perceived visual properties of a category. In the expert study, we found that both bird and dog experts listed more visual features for objects in their domain of expertise than novices. Closer examination of their features revealed that bird and dog experts were more knowledgeable about the color and shape properties of objects in their domain than novices. Given that bird watching and dog judging

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

require the ability to make fine perceptual discriminations quickly and accurately, it is not surprising that experts would be keenly aware of the visual properties that characterize objects in their domain of expertise. Thus, while the visual properties of an object are perceptually available to any observer, these features are not attended to unless they are relevant to the pragmatics of the observer's task.

The role of pragmatics can be further examined by studying experts who share the same object domain, but different task goals within that domain. Medin et al. (1997) investigated the category knowledge of three types of tree experts, taxonomists, landscape workers, and park maintenance personnel. They found that taxonomists and maintenance workers identified different morphological properties of trees whereas landscape workers differed from the other two types of experts with respect to the functional properties associated with trees. Although all three types of expertise required specialized knowledge of trees, the kind of knowledge needed to fulfill the pragmatics of that expertise differed across the groups. Other studies have shown that object properties are variable and individualistic even for experts who share the same object domain *and* task goals within that domain. Studying expert recreational fishermen, Boster and Johnson (1989) found that whereas some of the sport fishermen were sensitive to the morphological properties of fish, other sport fishermen placed more emphasis on their functional properties. These studies are informative because they show that despite sharing the same knowledge domain (i.e., trees or fish), there is considerable variation across different kinds of experts and also within the same kind of expert group.

In summary, the expert research demonstrates that the category properties of birds, dogs, trees, and fish cannot be determined independent of the person doing the categorizing. While these experiments have focused on expertise related to natural kind categories, there would be no reason to expect that the pragmatics of expertise would be any different for experts who specialize in the categorization of artifactual objects (e.g., automobiles, stamps, antiques). Finally, while the pragmatic view emphasizes the place of individual differences in object semantics, this approach does not deny that objects are predisposed toward certain visual and functional characteristics (as suggested by H&F's HIT model). However, beyond an object's normative properties, the manner in which an object is perceived and used will depend on the unique goals of the individual. Thus, the pragmatic view holds that object knowledge is not only determined by the structure found in the world, but also by the structure found in the mind of the categorizer.

ACKNOWLEDGMENT

This commentary was supported by funds from the National Science Foundation (SBR-9729030).

About numbers as a semantic category

Marc Thioux, Eva Turconi, Emanuelle Palmers, and Xavier Seron

Unité de Neuropsychologie Université Catholique de Louvain, 1348 Louvain-la-Neuve, Belgium. **{thioux; turconi; palmers; seron}@neco.ucl.ac.be**
www.neco.ucl.ac.be

Abstract: Humphreys & Forde concentrate on the living/nonliving dissociation. However, further dissociations have been reported, including selective loss or preservation in recognizing body parts and numbers. This commentary outlines the relevance of the number category for understanding the organising principles of semantic memory.

The observation of a double dissociation between numbers and other categories of words consecutively to brain lesions highlights the importance of numbers as a semantic category. Cipolotti et al. (1991) reported the case of CG, who could neither produce nor understand numbers, despite normal performance outside of the number domain. CG was totally unable to read aloud Arabic numerals, to write Arabic numerals to dictation, to discriminate Arabic numerals from meaningless shapes, to discriminate verbal numerals from nonwords, to produce orally the numeral that follows or precedes a given auditory numeral, and to judge which of two verbal numerals or which of two dot patterns is the larger. In contrast, she remained able to produce a normal amount of words from designed semantic categories, to name pictures from different semantic fields, to point to the picture corresponding to a spoken name, to define proverbs, to judge whether two words were synonyms, and her performance was also flawless in size judgement tasks and size serialisation tasks with pictures.

On the other hand, we have examined the case of a patient (NM) showing the exact reverse dissociation (Thioux et al. 1998). NM scored between 0 and 10% correct on tasks of picture naming, naming from a definition, and production of synonyms. Furthermore, NM's performance was also impaired in various comprehension tasks (identically for words and pictures and across different categories). Despite the presence of this semantic deficit and the high degree of anomia, NM retained preserved ability to comprehend and to produce numbers after the cerebral trauma. The dissociation was observed in oral naming, written naming, and in several tasks that do not require the processing of Arabic numerals (e.g., verbal fluency tasks, verbal association, recitation, "what comes next/before" questions). In addition, NM had retained preserved performance on tasks specific to the numerical domain. For instance, he could read aloud syntactically complex Arabic numerals, and he could compute or retrieve the solution of simple arithmetical problems.

One potential explanation of the preservation of numbers in NM was the relatively high frequency of these elements. However, we showed that the dissociation remained even when the items were matched for frequency. Another potential account of the preservation of numbers was the relatively small size of the category. However, there was no evidence of a strong correlation between the size of the category and NM's performance on tasks of naming or verbal fluency. Moreover, NM's performance was inaccurate in several very small categories composed of very frequent elements (e.g., seasons, geometrical figures). Thus, together with CG who presented with the reverse dissociation, the case of NM suggests that numerals form a relevant category at the semantic level. The only exception to the pattern of dissociation observed in these two patients was for the days of the week and the months of the year. In both patients, the series of days and months were either impaired or preserved with numbers.

Because recent neuroimaging studies have demonstrated that the left and right intraparietal sulci are sensitive to the magnitude of numbers (Dehaene et al. 1999; Pinnel et al. 1999), and because the left parietal region has been found to be the site of the lesion in several acalculic patients (e.g., Takayama et al. 1994), the selective deficit or preservation of the category of numbers can be attributed to the selective deficit or preservation of these cortical regions. Two different hypothesis have been advanced to explain why the semantic representation of numbers is dependent on the integrity of these specific cortical regions. According to the first hypothesis, the meaning of numbers is located where a prelinguistic system devoted to the estimation of numerosity has evolved (Dehaene 1997; Spelke & Dehaene 1999). This archaic system for processing numerosity is thought to emerge through the course of evolution within a population of cells involved in representing the location of objects in space independently from their identity and size (Dehaene & Changeux 1993). According to the second hypothesis, the meaning of numbers has come to be stored in a region responsible for visuo-spatial processing (Simon 1999). During child development, this non-numerical system would be solicited when learning to count (especially for finger-counting), and then for the construction of a representation of number magnitude.

Now, the question remains to know how many and what kind of

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

semantic attributes of numbers must be stored in the adult parietal cortex in order to explain the pattern of dissociation and association observed after a brain lesion. The current debate about category-specific impairments is mainly devoted to the question of the relative importance of visuo-semantic attributes in the processing of living and nonliving entities. If one turns towards numbers, it appears that, rather than being defined by visual or functional attributes, numbers are better defined by other numbers. In fact, numbers share a series of very specific semantic attributes: they pertain to the same ordered sequence, they represent a quantity, and this quantity can be defined as the sum of the preceding number and 1, when 0 and 1 have been defined (Russell 1921). Furthermore, numbers are subject to specific semantic manipulations such as subtraction, multiplication, parity judgment, and so on. In this respect, one might suggest that numbers are represented in an autonomous semantic store located in specific cortical areas; a true categorical account (TCA). In such perspective, most if not all, of the attributes of numbers must be stored in the same region, and all must be specific to numbers. However, in apparent contradiction with the TCA is the fact that numbers are also involved in the definition of other concepts (e.g., a dog has got four legs, there are 60 seconds in one minute, etc.). Another potential explanation of the category-specific deficit or preservation for numbers might be that they emerge as a category in the semantic system because they share some particular properties implemented in distinct brain regions; a reductionist account (RA). One such property could be that numbers are organised in an ordered sequence. At the semantic level, this means that each number must be linked to the following one by a "plus-one" link, and to the preceding one by a "minus-one" link (Thioux et al. 1998). The observation that, in both NM and CG, the categories of numbers, days, and months were conjointly spared or impaired fits with the hypothesis that these concepts have this important property in common. Nonetheless, a RA might have difficultly in explaining why different aspects of number meaning (e.g., order and magnitude) were conjointly spared or preserved in CG and NM. At first glance, this suggests that different meanings of numbers are grouped together in the same cortical areas.

In this respect, a RA can hardly be distinguished from a TCA, the only difference being that in the latter the critical areas are assumed to be involved in the storage of all the semantic attributes of the concepts, and only those. The fact that bilateral parietal areas are involved in eye movement, visuo-spatial representations, grasping, time estimation, as well as in magnitude estimation gives so many opportunities to justify one and the other account. In fact, all these abilities might have been important in the development of a fully integrated representation of number meaning, but also, each of these abilities might rely on a specific area within the parietal lobe.

Concepts and categories: What is the evidence for neural specialisation?

Lorraine K. Tyler and Helen E. Moss

Centre for Speech and Language, Department of Experimental Psychology, University of Cambridge, Cambridge CB2 3EB, United Kingdom.
lktyler@csl.psychol.cam.ac.uk **hem10@cam.ac.uk**

Abstract: Humphreys and Forde argue that semantic memory is divided into separate substores for different kinds of information. However, the neuro-imaging results cited in support of this view are inconsistent and often methodologically and statistically unreliable. Our own data indicate no regional specialisation as a function of semantic category or domain and support instead a distributed unitary account.

Humphreys & Forde's (H&F's)Hierarchical Interactive Theory (HIT) of object recognition makes important claims about the organisation and processing of semantic knowledge. In this commentary, we concentrate on two main issues: (1) the authors' interpretation of the neuro-imaging data as support for their claim that semantic knowledge is divided into separate stores and that there is top-down activation of visual information in naming, and (2) the relation between the organisation of semantic knowledge and the internal structure of concepts.

H&F suggest that "the concept of a semantic system in any unitary sense may be one of the victims of an attempt to define the nature of our stored knowledge in more detail" (p. 6). They argue for multiple substores, concentrating on the separation between stored structural descriptions and "stored functional and inter-object association information." Much of the motivation here comes from the analysis of the neuroimaging literature (see sect. 5.1) where H&F interpret existing studies as supporting the claim for anatomical differences in the representation of different categories or types of knowledge. However, close inspection of the neuroimaging data suggests a more complex and less coherent picture than is implied in their review. Instead of a consistent set of regions being activated in response to either living or nonliving things, across studies many different regions are activated. For example, Perani et al. (1995) report activation in the left fusiform and the left lingual gyrus for natural kind concepts. None of these regions overlap with those activated in Mummery et al. (1996) which reports activation in the left anterior temporal pole, right anterior temporal pole and right inferior parietal lobe. This discrepancy is typical across the board, with perhaps the exception of the left posterior middle temporal gyrus which has been activated in a number of studies in response to tools (Damasio et al. 1996; Martin et al. 1996). However, even this region is not robustly activated within individual studies.

This general lack of consensus across studies is due to many factors, some methodological and some statistical. H&F acknowledge that some differences may "reflect variations in scanning procedures or in the task requirements." We believe that the problems are even more wide-ranging than this, and warrant greater caution in the interpretation of results. For example, sets of stimuli are not always controlled for factors such as familiarity, imageability, and letter length (for words) and visual complexity or structural similarity (for pictures) although these variables can generate spurious differences (e.g., Funnell & Sheridan 1992). Problems also arise because in many studies statistical maps are not corrected for multiple comparisons between voxels. Using uncorrected alpha levels over large numbers of contrasts increases the probability of obtaining a high proportion of false positives. Thus, many of the reported significant activations may, in fact, be false positives. Indeed, some studies include little statistical support for apparent differential activations. Chao et al. (1999) for example, do not report full statistical information, such as z scores and their accompanying significance values (nor do they correct for multiple contrasts), yet H&F interpret their finding of lateral fusiform activation for questions about living things as strong support for their claim that specific forms of knowledge are drawn on to different degrees when accessing knowledge about living and nonliving things.

When we try to avoid these problems by controlling for nuisance variables and correcting for multiple statistical comparisons, we find that there is no neural specialisation as a function of category or domain of knowledge (Devlin et al., in press; Tyler & Tovares, in press). In three PET studies using different tasks (lexical decision and semantic categorisation) we found an extensive network in L fronto-temporal regions which was reliably activated for semantic processing of words and pictures. However, there were no domain or category effects in either study. We only found differences between natural kinds and artefacts at an uncorrected level of significance and these activations were generally inconsistent across the experiments.

H&F also rely on neuroimaging results to support the idea of re-entrant activation of visual information in the naming task. Their argument is as follows: (1) left inferior and posterior temporal lobes are more involved in object naming than in naming the

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

colours of non-objects (Price et al. 1996), (2) if inferior and posterior temporal regions are linked to stored visual knowledge, then this shows that naming known objects involves extra recruitment of visual knowledge over and above simply recognising those objects (i.e., re-entrant activation). The basis for assuming that these regions are linked to stored visual knowledge is that they are activated to a greater extent for animals than tools. However, as we have argued above, the support for this claim is inconsistent at best. Moreover, even if these regions are more activated for animals than tools, it seems circular to assume that this is because of the greater demand on visual processing for these concepts. Thus, although we agree that interaction among the different stages of the hierarchy may be more plausible than the strictly feed-forward cascade model, the specific support cited for this claim is not clear cut.

Thus, the neuroimaging data do not provide compelling support for neural specialisation as a function of either category or domain. Indeed, they appear to be more compatible with the kind of distributed account of conceptual knowledge which we have been developing (Durrant-Peatfield et al. 1997; Tyler & Moss, in press; Tyler et al. 2000). In this account, concepts are represented as sets of overlapping features – some distinctive and some shared across concepts. Frequently co-occurring features (has-legs, has-ears) support each other and are thus more resistant to damage. Since their distribution differs across categories, some categories are more susceptible to semantic damage than others, leading to category-specific semantic impairments (Moss & Tyler 2000; Tyler et al. 2000). On this account, there is no anatomical differentiation by semantic content. Concepts are represented in a distributed neural system including frontal, temporal, parietal, and occipital regions. The components of this highly interactive system will be more or less involved, depending on a number of factors – the nature of the input, specific task, and the additional nonlinguistic cognitive demands required.

This leads to our second point. H&F discuss our account of conceptual structure, and agree that "the differences between shared and distinctive features and the degree to which these features correlated with the function of the object are likely to be important contributing factors in category-specific deficits." We acknowledge that such models need to be elaborated further and set within a framework for object recognition (and language comprehension). In fact, we have always assumed (implicitly, at least) that there is a structural description system (or some form of low-level object processing) necessary for object recognition. We differ from H&F in assuming that perceptual properties of objects are not only represented at this level – in a modality specific store of "templates" for object recognition – but that they are also represented along with other kinds of semantic property within a distributed semantic system. One reason for this is that structural descriptions seem to consist of only a subset of the perceptual properties of an object – essentially its shape – leaving other perceptual properties such as colour, size, surface texture, and pattern unaccounted for. Second, how are correlations among the perceptual and functional properties of objects captured if these properties are stored independently in modular subsystems? H&F appear to be in agreement with an approach that emphasises the internal structure of concepts – yet it is not clear how they are proposing these inter-feature relations are to be instantiated across the different information stores.

Thus, while disagreeing with some of their conclusions, we support H&F's efforts in laying the groundwork for a thorough discussion of the issues involved in constructing a cognitive and neural account of object recognition.

Putting semantics back into the semantic representation of living things

Deborah Zaitchik[a] and Gregg E. A. Solomon[b]

[a]*Department of Psychiatry, Massachusetts General Hospital, Boston, MA 02114;* [b]*Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139.*

zaitchik@psych.mgh.harvard.edu **geas@psyche.mit.edu**

Abstract: The authors' model reduces the literature on conceptual representation to a single node: "encyclopedic knowledge." The structure of conceptual knowledge is not so trivial. By ignoring the phenomena central to reasoning about living things, the authors base their dismissal of semantic systems on inadequate descriptive ground. A better descriptive account is available in the conceptual development literature. Neuropsychologists could import the insights and tasks from cognitive development to improve their studies.

The authors claim to be modeling semantic representations in the domain of living things. To the extent that semantics is about meaning, we argue that the authors have captured only a small piece of the puzzle. This point is agnostic with respect to the potential adequacy of a PDP account of the phenomena. Rather, it is a point about what any model will have to take into account if it hopes to characterize this cognitive domain adequately. Briefly, what we do not see in the present paper, nor in most of the neuroscience literature at large, is the link between semantics and concepts.

The sense from much of the neuropsychology literature is that semantics does not go much beyond naming and object recognition. However, underlying the semantic phenomena measured in such tasks is a rich conceptual structure. While this structure is acknowledged in passing by the authors, it plays no role in the model; "encyclopedic knowledge" is a single node, as if its retrieval were a trivial matter, not one involving rich and variegated inferences. For this model, then, a semantic system is not necessary, but only because it assumes successful retrieval of the very information such a system would provide.

In fact, the domain of living things is especially rich (see Medin & Atran 1999). It is one of the earliest and largest domains of knowledge represented by children; we may even have innate perceptual and attentional mechanisms for the detection of living things (such as specific motion detectors). Thus, it has been extensively studied by cognitive developmentalists. This literature provides us with a rich description of the central phenomena that characterize reasoning about living things. The best account of these phenomena to date is that semantic representations in this domain are organized in a hierarchical conceptual structure. Moreover, we represent certain concepts as having special status within the domain. Biological concepts such as growth, nutrition, movement, and reproduction play a central role in how we reason about living things, allowing us to make predictions and supporting causal explanations. Core concepts such as these specify which objects are in the domain and which are out.

To the extent that researchers want to model semantic representations of living things, this is the understanding that needs to be explained. It is an understanding that is not captured by a model of encyclopedic knowledge that likens retrieval to looking up meaning in a dictionary. For example, there is evidence that children's understanding of what is "alive" undergoes radical conceptual change – a change more fundamental than simply altering a denotational listing (Carey 1985). We understand the meaning of the concept "alive" in terms of a system of interrelated concepts. For children to understand what is alive as we adults do (i.e., for them to make similar predictions or explanations about the behavior and attributes of living things), they must come to interrelate the various phenomena of living things in terms of the causal mechanisms that govern the bodily machine. Moreover, there are complex and nontrivial ways in which children's understandings diverge from those of adults even when they know the facts relevant to the adult understanding, as is the case in their rea-

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

soning about biological inheritance and germ-based illness (Solomon et al. 1996). There are complex and nontrivial paths by which children come to reason about living things as do adults in their culture (Hatano & Inagaki 1994). And there are complex and nontrivial ways in which reasoning about living things breaks down, in which the node of encyclopedic knowledge fails to provide the relevant conceptual information, as we see in other studies of patients with Alzheimer's disease (Chan et al. 1997).

Encyclopedic knowledge is the heart of semantics; characterizing it as a single node, and then largely ignoring it, simply will not do. If the goal of the endeavor is to capture everyday reasoning and changes in it with time, development, or impairment, then the model must be extended so as to take account of conceptual structure. To be sure, this will require the use of a new set of tasks. Though the standard neuropsychological tasks (e.g., naming, picture recognition) were critical in revealing semantic impairment, they cannot lay bare the underlying conceptual domain. Thus they cannot address the variety of ways in which the concepts of living things can be acquired and can break down. The authors, like others in the field, make broad claims about their ability to model our reasoning about living things while looking at only a tiny fraction of the data.

Researchers in cognitive development have laid out a comprehensive picture of the semantic relations normally represented in this domain from infancy to adulthood. They have a large and well-established literature utilizing a wide array of well-understood methods. To date there has been little cross-fertilization between cognitive developmentalists and neuropsychologists (the work of Chan is a notable exception). By importing the insights and methods from research in conceptual development, neuropsychologists can better address the debate about cognitive architecture. We worry that the success of the authors' model, as it stands, is due to its limited notion of semantic knowledge. It is possible that such PDP models could be extended successfully to account for reasoning about living things, but without a more complete descriptive account of the conceptual domain, we cannot say that the debate has yet been adequately engaged. We are not endorsing PDP; we are endorsing the debate.

Authors' Response

Category specificity in mind and brain?

Glyn W. Humphreys[a] and Emer M. E. Forde[b]

[a]*Behavioural Brain Sciences, School of Psychology, University of Birmingham, Birmingham B15 2TT, United Kingdom;* [b]*Institute of Psychology, Aston University, Aston, Birmingham, United Kingdom.*

g.w.humphreys@bham.ac.uk **e.m.e.forde@aston.ac.uk**

Abstract: We summarise and respond to the main points made by the commentators on our target article, which concern: (1) whether structural similarity can play a causal role in normal object identification and in neuropsychological deficits for living things, (2) the nature of our structural knowledge of the world, (3) the relations between sensory and functional knowledge of objects, and the nature of our functional knowledge about living things, (4) whether we need to posit a "core" semantic system, (5) arguments that can be marshalled from evidence on functional imaging, (6) the causal mechanisms by which category differences can emerge in object representations, and (7) the nature of our knowledge about categories other than living and nonliving things. We also highlight points raised in our article that seem to be accepted.

R1. Sometimes it's not what you say but what you don't say

Sometimes it's not what people say but what they don't say that makes the interesting point. Our target article, on the nature of apparent category-specific deficits in neuropsychological populations, had several objectives. These were (1) to review both the empirical database and the different theoretical accounts of category-specific deficits within the literature; (2) to argue that no single account was able to explain the full pattern of deficits that have been reported; and (3) to propose that a deficit at one locus of processing could not explain all cases, and that patients could have impairments at various stages of processing (including deficits that are more "perceptual" than "semantic" in nature), each of which could generate a deficit for particular categories of object. Finally, we attempted to provide an overall processing architecture, the Hierarchical Interactive Theory (HIT) to account for the different deficits that have been documented. The first main assumption of HIT is that the stages of processing leading to object naming are hierarchical, with (for example) stored perceptual knowledge being distinct from, and accessed prior to, stored functional knowledge about objects. The second main assumption is that the object recognition system is highly interactive, with re-entrant feed-back activation as well as feed-forward activation between the different stages of processing.

The commentators on our article raise many interesting points, summarised in Table R1. Only **Bub & Bukach** comment (favourably) on our argument concerning the heterogeneity of the disorders, none comment on the evidence for a hierarchical organisation of the object processing system, and only **Pulvermüller** discusses the idea of re-entrant feedback in any detail. This silence is important. From it we infer that there are indeed no major objections to our contentions on these points; the commentators accept that the functional deficit can vary across patients (and can include forms of perceptual disorder), and that perceptual knowledge is distinct (and can be dissociated from) other forms of stored knowledge about objects. From this it follows that category-specific deficits should not be treated as a unitary phenomenon, rather the functional locus of the problem should be defined in each patient and studies need to clarify the particular type of patient that is being accounted for. This is important because in previous investigations researchers have assumed that, because an impairment may be isolated to the same category of object in different patients, the locus of the deficit is the same (typically being thought to derive from a disorder of semantic memory; e.g., Caramazza 1998). Our argument is that this assumption is incorrect. If this point is taken on board in the literature, then we believe that our target article has been worthwhile, even if the full details of our own account are not accepted or need further working through. It is the particular details of our account that commentators have in the main responded to, and which we answer below.

The major issues that are highlighted by the commentators are listed in Table R1, along with the commentators involved.

R2. Structural similarity and category effects in normal object identification

In our target article, we proposed that, defined across a category, living things tend to have more similar perceptual

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

Table R1. *Issues raised by the commentators on our target article*

Issues:						
Structural similarity	Perceptual knowledge	Sensory-function relations	Semantic knowledge	Imaging	Underlying mechanisms	Other categories
Commentators:						
Laws	Arguin & Saumier	Cree & McRae	Bub & Bukach	Bub & Bukach	Arterberry	Thioux et al.
Leek & Pothos	Arterberry	Lambon Ralph & Garrard	Coslett & Saffran	Coslett & Saffran	Barrett	Pulvermuller
Mahon & Caramazza	Davidoff	Leek & Pathos	de Almeida	Mahon & Caramazza	Lamberts	
	Graf & Schneider	Mahon & Caramazza	Pillon & Samson	Tyler & Moss	Lambon-Ralph & Garrard	
	Mahon & Caramazza	Pillon & Samson	Tyler & Moss		Tanaka	
	Saffran & Coslett		Zaitchik & Solomon		Zaitchik & Solomon	

structures than nonliving things (one exception being body parts, which are sometimes classified as "living things" and which tend to be perceptually dissimilar within their category; see Humphreys et al. 1988). Two sets of commentators, **Mahon & Caramazza** and **Laws**, query whether this is the case. Mahon & Caramazza criticize a paper by Humphreys et al. (1995), which used similarity ratings on subsets of living and nonliving things to define the input representations for simulations of category-effects in both normal and impaired object identification. They argue that ratings generated across subsets of objects may be biased by the objects chosen. Thus, not too much should be made of the fact that participants rated living things as being more perceptually similar than nonliving things with the sample of objects used in that paper. We agree with this point. However, we are puzzled by the reference to this one paper which used a subset of objects for simulation purposes, when, in the target article, we mentioned other studies that have attempted to generate independent measures of perceptual similarity within theoretically-unbiased samples of object categories. For example, Humphreys et al. (1988) reported results from measures of within-category overlap across the bounding contours of 76 name-frequency matched objects depicted in the Snodgrass and Vanderwart (1980) picture norms. Measured across all the objects depicted within the different categories, living things had higher levels of contour overlap than nonliving things (barring body parts). Humphreys et al. also had independent participants list the number of parts present in objects and reported that living things tended to have more common parts than nonliving things. These other studies provide independent support for the argument about differences in perceptual similarity within classes of living and nonliving things.

Laws raises a related point to **Mahon & Caramazza**. In fact, he argues that nonliving things are structurally more similar than living things, if similarity is based on the number of pixels in common across the exemplars in the Snodgrass and Vanderwart norms. He suggests that this provides a more valid measure than contour overlap because it is sensitive to the internal detail in the depictions. The trouble here is that the overlap in the internal detail specified in line drawings is quite arbitrary and likely to reflect whether the artist has drawn in a texture. Pixel-level measures, then, are perhaps less likely than outline contour similarity to reflect the perceptual relations between objects in real life. Indeed, measures of contour overlap have been shown to predict the degree of naming deficit in at least some patients with apparent category-specific deficits for living things (Forde et al. 1997). We also note that other attempts to develop pixel-level measures of similarity do not agree with Laws's conclusion. Gale et al. (in press) have simulated category-specific deficits in object recognition in which a first stage of object classification was carried out by means of a Kohonen map (Kohonen 1988). They used pixel-based input descriptions of photographs drawn from sets of animals, furniture, musical instruments, and clothing (35 per category). After training the model they found that, within the Kohonen map, animals and musical instruments tended to have closely clustered representations, with individual objects being represented by overlapping units each activated to a moderate degree. In contrast, individual items of furniture and clothing tended to be represented by fewer units with sharp peaks of activation. Identification in the model was assessed by training a further module to assign names to the individual objects. Gale et al. found that it was more difficult for the model to learn basic names for animals and musical instruments than for furniture and clothing, reflecting the overlap in the representations for animals and musical instruments within the Kohonen map. The interesting point about this simulation is that the categorical differences emerged despite the model not being provided with any explicit information about object categories. This provides independent, theoretically-neutral evidence for similarity being higher for at least one major class of living thing (animals) than for two

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

classes of nonliving thing (furniture and clothing). It is also interesting that, for the model, musical instruments tended to be treated more like animals than to the other nonliving things. In the neuropsychological literature patients with problems in identifying living things often also have problems with musical instruments (e.g., Warrington & Shallice 1984; see also **Leek & Pothos**). Gale et al.'s simulations suggest that this is due to inherent differences in perceptual similarity across these object classes.

An experimental approach to the same issue has also been reported by Lamberts and Shapiro (in press). They had normal participants perform a task requiring access only to perceptual information about the stimuli – simple same-different matches with living and nonliving stimuli, with items on "different" trials being paired at random within each category. Living things were matched more slowly than nonliving things, consistent with them having higher levels of perceived similarity.

Laws also queries whether normal participants show an advantage for identifying nonliving over living things, as might be expected if living things have greater levels of perceptual similarity across category members. There are a number of points to be made here. One is that the argument is based on studies by Laws and colleagues with a relatively small set of items which tended to be low in familiarity and name frequency. In other experiments measuring the speed of normal object naming, perceptual similarity has been found to vary with name frequency, with the largest effects being found on items with high frequency names (Humphreys et al. 1988). We have argued that this is because of competition from perceptual neighbours for items from categories with high perceptual similarity, which counteracts the advantage for objects with high frequency names. Hence studies using low frequency items are unlikely to be sensitive to the most telling effects of perceptual similarity. In the data of Laws (1999) category was also found to interact with gender – males tending to show an advantage for nonliving things and females for living. Again this may reflect the low familiarity of some of their nonliving things for the female subjects. In addition, on our account it is perceptual similarity, rather than category, which should be important for normal participants. Laws and colleagues have only varied category. Humphreys et al. varied perceptual similarity and so included body parts within their perceptually dissimilar items. Manipulations of category alone may cut across some variations in perceptual similarity. Laws et al. further suggest that the data from Humphreys et al. could have been due to variations in factors such as the familiarity and complexity of the items. However, when these data are re-analysed with familiarity, complexity, and other factors such as name and image agreement and individual name frequency as co-variates, then both the effect of category and the category x frequency interaction remain ($F(1,67) = 9.51$ and $F(1,67) = 6.50$, both $p < 0.025$). One other study reporting an advantage for structurally dissimilar over similar living things in normal object naming was reported by Snodgrass and Yuditsky (1996), using a considerably larger set of items from the Snodgrass and Vanderwart norms than any other study. Again re-analysis of the data across all categories with structurally similar and dissimilar items (taking the categories from Humphreys et al. 1988 – a set of 216 items in total) reveals that the category effect remains even with name and image agreement, familiarity, complexity, and individual name frequencies filtered out ($F(1,209) = 4.51$, $p < 0.04$). These prior results do not seem to have been due to the effects of covarying factors.

The results of the studies by **Laws** and colleagues nevertheless indicate that it is important to understand the circumstances in which apparent category effects can occur for normal participants. In some of the studies by Laws (e.g., Laws & Neve 1999), living things were better identified than nonliving things under brief presentation conditions. Gerlach (in press) has replicated this advantage for living things when objects are degraded by lateralised stimulus presentations. However, using exactly the same items, Gerlach found the opposite advantage (for nonliving things) under unlimited presentation conditions. Gerlach used an object decision task and reported that the nonliving advantage increased as the specificity of the information required for object decision increased (e.g., as the non-objects became more object-like). This last result is consistent with perceptual similarity having an impact when observers conducted a detailed assessment of stored visual knowledge to discriminate objects from non-objects under normal viewing conditions. Under degraded presentation conditions, though, several factors may act against nonliving things including: the loss of high spatial frequency components, increased variation in base-level exemplars which makes their perceptual structures less reliable when degraded (see Laws & Neve 1999, for this argument), and the dependence of living things on global shape properties for their identification. At present it is unclear which (if any) of these factors is crucial, though it is clearly important to understand which perceptual factors determine the efficiency of identification, and under what conditions, if we are to derive a full account of why effects of object category emerge even with normal observers. Nevertheless, we propose that important forms of perceptual similarity are high amongst many categories of living thing, and that perceptual similarity has a deleterious effect on normal object identification under standard viewing conditions. This is further supported by evidence that some patients with apparent category-specific deficits for living things also fare badly when asked to name perceptually similar sets of nonliving things (Forde et al. 1997); by evidence on the clustering of visual and semantic errors when normal subjects make naming errors to objects (Vitkovitch et al. 1993); and by the evidence on "category differences" emerging when nonhuman primates have to learn individual responses to living and nonliving things (Gaffan & Heywood 1993).

R2.1. The nature of our perceptual knowledge and the necessity and sufficiency of access to different knowledge representations

In our target article, we distinguished between stored structural knowledge about objects (which we termed our stored structural descriptions) and our knowledge about inter-object associations and object function. This distinction was derived from neuropsychological data indicating that patients could have intact visual access to structural knowledge (as indicated by good performance on difficult object decision tasks), as well as impaired access to semantic knowledge (judged by poor ability to match associatively related objects). As we indicated in the article (n. 3), we are not committed to our structural knowledge of objects taking a particular form of representation – for ex-

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

ample, this knowledge could explicitly represent the parts of objects and their inter-relations or it could be more holistic, without explicit part decomposition. Although the term "structural description" can imply a parts-based representation (e.g., Biederman 1987), in our hands this is not the case, and we used the term primarily to maintain consistency with earlier articles in which the distinction between perceptual and other forms of stored knowledge was derived (e.g., Humphreys et al. 1988). **Graf & Schneider** discuss whether structural knowledge is holistic or parts-based. They are quite correct to point out that evidence on the limits of view-invariance in object recognition goes against many part-based accounts of structural knowledge, though it may be possible to salvage a parts-based account using implicit coding, as they suggest. However this work on view invariance goes beyond the scope of our target article.

Nevertheless, arguments about the nature of our stored perceptual knowledge do have an impact on the debate concerning category-specific deficits. Let us begin by considering knowledge of object form, before discussing the representation of other perceptual attributes (colour, size, etc.). As far as form information is concerned, it may be that stored knowledge differs in kind for contrasting object categories. For example, parts-based decomposition may be more important for the recognition of nonliving things than the recognition of living things, and that representations that code the relative positions of parts may mediate the recognition of nonliving but not living things. For animate objects, at least, parts may move with respect to one another, making a representation dependent on relative position coding unreliable. On the other hand, for inanimate objects the spatial relations between the parts can dictate the functions of the stimuli (e.g., a handle in the correct relation to a blade), and so may be strongly weighted in the identification process. **Davidoff** mentions some data compatible with this. He noted evidence from Davidoff and Warrington (1999) on a patient who could name animals but could not discriminate their parts, consistent with a parts-based description not being necessary for these stimuli. In order to advance the debate on category-specificity further, theories need to specify in more detail the nature of the perceptual representations for different object types.

Arguin & Saumier make a somewhat different point. They discuss the arguments in our article concerning the role of action-based knowledge in the identification of nonliving things, and suggest that "the role of action-related functional knowledge in the processing of nonliving things is directly analogous to that of visual structural knowledge . . . (in) the recognition of living things." If this is a debate about the relative "weighting" of action and structural knowledge in the identification process, then Arguin & Saumier's statement does capture our point. We argue that action knowledge is strongly weighted for the identification of nonliving things. Nevertheless, we continue to think that, in the main part, access to structural knowledge is *necessary* in order to identify both types of stimulus[1] and, contrary to Arguin & Saumier's inference, we hold that there are representations of structural knowledge for nonliving things. It is simply that structural representations may differ for different stimuli (see above), and that structural knowledge plays a lesser role in determining the efficiency of object identification for nonliving things than for living things.

Arterberry in fact criticises our distinction between structural knowledge of nonliving things and knowledge about their functions (e.g., how such objects are used). She points out that some theorists maintain that "visual/perceptual information cannot be dissociated from function (or action) information" (e.g., Gibson 1979). Our take on this is slightly different. We propose that perceptual knowledge can be directly linked to associated body actions independent of contextual knowledge or knowledge about inter-object associations, as highlighted by **Saffran & Coslett** in their commentary. Such perception-action linkages will operate for nonliving things in particular, which serve as the implements for our actions, rather than for living things which are animate in their own right.

However, although we argue that there is differential weighting across categories for various forms of knowledge about objects, the jury is open on whether access to particular forms of knowledge is *sufficient* to enable identification to occur. In his commentary, **Davidoff** interprets our argument as being that "only sensory knowledge . . . has a strong (direct) access to the phonological lexicon," as if access to perceptual knowledge were sufficient for object naming. This possibility is also discussed by **Pulvermüller** as one of the "paths through the network" that may exist. However, we are not advocating a direct path from perceptual knowledge to names. To the best of our knowledge, there is no strong evidence to indicate that access to sensory/perceptual knowledge is sufficient for the naming of living things (for example), in the sense that patients can name these items without retrieving any other form of knowledge. We are also not convinced that there are direct links between sensory properties such as colour and object names (see Humphreys & Price 2001). It seems more that sensory knowledge must combine with other forms of functional and associative knowledge to enable name retrieval to take place. But within this combination process, sensory/perceptual knowledge is weighted strongly for living things.

One other relevant point here concerns the relations between structural and other forms of perceptual knowledge about objects, such as their colour. Again, we were relatively neutral on this issue in our target article, though we think it is likely that structural knowledge about objects is represented independently of knowledge about (for example) colour – as evidenced by patients who show good access to structural knowledge along with impaired access to coloured knowledge about objects (see Price & Humphreys 1989 for discussion), or impaired colour knowledge without problems in accessing shape knowledge for object identification (Luzzatti & Davidoff 1994; see **Davidoff**). This also relates to the debate about which properties are necessary and which are sufficient for object identification to take place. Within a system with separate representations of different perceptual attributes, identification may be contingent on access to some but not other forms of representation. We suggest that identification does depend in large part on access to stored structural knowledge about objects, since shape is diagnostic of object identity (though not necessarily uniquely diagnostic, see n. 1). Other forms of perceptual representation may not be necessary (indeed, we can recognise objects from line drawings with no colour information present), but they may still contribute. Access to forms of functional knowledge, on the other hand, may be necessary in addition to access to stored structural knowledge.

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

We have also argued that damage to visual knowledge about objects can differentially impair the recognition of living things. This can be one cause of an apparent category-specific impairment. The evidence favouring this comes from studies of patients with perceptual rather than semantic impairments in object processing (e.g., Riddoch & Humphreys 1987) and from evidence that deficits for living things can be enhanced by increasing the difficulty of perceptual processing (Humphreys et al. 1999). Tangental to this point, **Mahon & Caramazza** argue that visual deficits cannot necessarily cause difficulties with living things because there are patients with deficits in visual knowledge without having greater deficits with living things. In our view, the patient data on this point are not particularly strong and one needs to be cautious in making inferences from small databases. Nevertheless, we discussed the point in our target article, where we summarised simulations that showed that either differential or similar deficits could emerge for living and nonliving things depending on how damage was inflicted on visual representations. In simulations with randomly fluctuating noise, for instance, we found that deficits for living things were not significantly greater than for nonliving things. Such simulations show that, at least in principle, category-specific deficits are not a necessary consequence of visual damage. The nature of the visual/perceptual impairment may also be important if the role of parts versus overall shape differs across categories (see above). Nevertheless it is important to reiterate that these points do *not* rule out the argument that impairments in visual processing and representation can be sufficient to generate these deficits. Necessity and sufficiency should not be confused here.

R2.2. Sensory-functional relations

A further argument made in our article was that category-specific deficits could also be determined by contrasting patterns of correlation between the sensory/perceptual properties of objects and knowledge about object functions, and also by there being different types of functional knowledge that characterise contrasting object types (e.g., knowledge of how to perform actions with objects vs. knowledge of biological functions). Several authors respond to this point, some in agreement and others who disagree that different patterns of correlation underlie some forms of category-specific deficit. **Mahon & Caramazza** argue that subjects list visual and functional/associative properties to the same degree when asked to generate property norms, inferring that these properties are weighted equally for the different object types. On the other hand **Cree & McRae** and **Lambon Ralph & Garrard** both argue that normative feature listings do vary across object categories – for example, the number of distinctive functional features being greater for artefacts than for living things. In addition, Lambon Ralph & Garrard point out that shared features are more prominent for living things than for nonliving things. Both differences may contribute to category-specific deficits emerging in patients. Cree & McRae and **Pillon & Samson** further note that some attributes may be more central to the representation of some objects than others, and this too may be important. One of the difficulties here, raised by Cree & McRae, is that a simple separation into sensory and functional features fails to pick up on finer-level variations in both types of knowledge that may prove to be important. We tried to make essentially this point in our target article when discussing the contrasting forms of "functional" knowledge that are often lumped together in studies. We suggested that contrasting types of functional knowledge may be differentially involved in the recognition of living and nonliving things. However, this is difficult to gauge from existing data because authors have used the term "functional knowledge" (and indeed the term "sensory knowledge") to refer to widely different types of information. This is highlighted by the different properties used to probe "sensory" and "functional" knowledge in relevant papers in the literature (see Table R2). In general, we disagree with those commentators who seemed to conclude that we were presenting a "standard" view that the distinction between sensory and functional features in general will be sufficient to account for all patients (**Mahon & Caramazza, Leek & Pothos**). Our view on this is more in line with the argument made by **Cree & McRae**, who suggest that a sub-division into eight knowledge types is needed to capture the variance in normative listings. Irrespective of the validity of this particular proposal, we believe that attempts to provide finer-grained characterisations of our knowledge about objects will help to give a fuller account of the neuropsychological data. Does such a finer-grained breakdown mean that our knowledge is domain-specific, as suggested by Mahon & Caramazza? We say no. Note that, in Cree & McRae's analysis of eight knowledge types, musical instruments clustered with creatures and nonliving food stuffs with fruit and vegetables. This clustering is dependent on detailed analysis of different knowledge types rather than object category. Note also that, in HIT, any differences in knowledge representation are linked to proposed variations in processing efficiency which are not domain-specific (e.g., due to differences in perceptual similarity). We hold that both differential weighting of knowledge types, and variations in processing as a function of different forms of similarity, are important in explaining the range of category effects reported in both normality and pathology, and that fine-grained category-specific impairments arise naturally from the view (expressed in HIT) that categories have different weightings across a range of knowledge stores.

In fact it may be that fine-grained breakdowns are more difficult for views such as the "domain specific theory" of semantic knowledge (e.g., Caramazza & Shelton 1998). On this account, more domain-specific systems would need to be posited to explain increasingly finer-grained deficits. In Caramazza and Shelton (1998), for example, three domains of knowledge were specified: animals, plants, and man-made objects. However, following the report of a patient with a selective sparing of the category body parts (along with impaired naming of other categories of object), Shelton et al. (1998) suggested that body parts formed a fourth domain-specific category of knowledge. Warrington and McCarthy (1987) reported an even finer-grained distinction between manipulable and nonmanipulable objects. Following the domain-specific account we might presume that there are separate systems for the recognition of these two classes of object, and so forth. This seems unparsimonious and unconstrained.

Mahon & Caramazza also assert that the main assumption of the sensory-functional account of category deficits has been refuted by patients who show impairments for living things but equivalent levels of performance with "func-

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

Table R2. *Properties used in tests of the "sensory-functional" account of category-specific disorders*

	Living things	Non-Living things
Sensory knowledge		
Caramazza & Shelton (1998)	constituent parts (mouth)	constituent parts (feathers)
		colour
		overall shape
		size
Lambon Ralph et al. (1998)	size	constituent parts (wheels?)
	colour	colour
	sound	tactile information
	tactile information	
	temperament (fierce?)	
Samson et al. (1998)	constituent parts	constituent parts (mast?)
	overall shape	overall shape
Functional Knowledge		
Caramazza & Shelton (1998)	biological functions (breathes?)	use for humans (protects?)
	habitat (live in water?)	
	associated action (flies?)	
	other (buy frozen?)	
Lambon Ralph et al. (1998)	habitat	physical properties (glass?)
	use for humans	use for humans
	associated action	used by whom (sailors?)
	sound made	context of use
	attributes of temperament (fierce)	
	food source (eats cheese?)	
Samson et al. (1998)	habitat	place where used (urban?)
	human use (transport?)	use for humans
	associated action	context of use
	taste	
	food source	

tional" and "sensory" questions. A problem here, though, again concerns the different ways in which the terms "functional" and "sensory" knowledge are operationalised across different papers (see Table R2), and even the different stimuli classed as living or nonliving (e.g., animals vs. plants and objects, Caramazza & Shelton 1998; animals vs. manmade objects, Lambon Ralph et al. 1998; animals and plants vs. man-made objects, Samson et al. 1998). For HIT, the different kinds of sensory and functional properties may be stored in different systems (colour knowledge, knowledge of shape, encyclopaedic knowledge, knowledge of characteristic action, etc.), so that tests of different properties should not be thought to address a common knowledge store.

R2.3. Do we need semantic representations?

Another issue discussed in our target article is whether we need to posit a core semantic system that "processes, stored, and retrieves information about the meaning of words, concepts, and facts," distinct from forms of perceptual and motor knowledge about objects. We speculated that this concept of core semantic knowledge may be a victim of the more detailed specification of knowledge types that we are advocating. On this view, "semantic memory" emerges from the inter-connections between different perceptual and action-based representations, perhaps augmented by encyclopaedic knowledge held in verbal long-term memory. This proposal is supported in the commentary of **Coslett & Saffran**.

Several other commentators argue differently and hold that the notion of a core semantic system should be maintained. **Tyler & Moss**, for example, suggest that perceptual knowledge is represented not only within modality-specific perceptual recognition systems (including a structural description system for object forms), but also within semantic memory. They ask "how are correlations among the perceptual and functional properties of objects captured if these properties are stored independently in modular subsystems?" **Pillon & Samson** similarly argue that a semantic system, separated from perceptual and motor knowledge, is needed to integrate information across the more peripheral knowledge stores. **Bub & Bukach** suggest that contextual contiguity may be important in our knowledge representations, and we suppose that this too may rely on some form of (semantic) abstraction beyond perceptual knowledge. **Zaitchik & Solomon** argue that semantic representations will need to be incorporated into theories when processes involved in reasoning must be taken into account, and **de Almeida** proposes that semantic representations are not decomposable into more primitive features (shape attributes, colours, smells), and hence semantic representations may be distinct from other forms of (decomposable) knowledge. These arguments all have some weight, but they are also difficult to resolve empiri-

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

cally. One approach, which we are sympathetic to, would be to start with the presumption that all our knowledge is either perceptuo-motor in nature or verbalised encyclopaedic information, and to see how far such an account can stand up to the empirical results. For example, in detailed simulations of this approach, can links be established between the different knowledge systems that incorporate our "semantic" knowledge about perceptual-action relations, contextual relations between objects and so forth? The relatively small-scale simulations that have been conducted to date are probably insufficient to know if other (core semantic) forms of representation are needed (e.g., Devlin et al. 1998; Durrant-Peatfield et al. 1997). Within the distributed approach we are suggesting (with no core semantic system), our notion is that the different knowledge systems are modular in the sense that they can be selectively damaged, but they are also highly interactive – for example, perceptual knowledge stores may need to be consulted to retrieve knowledge about the biological functions of animals. Hence, unlike **Tyler & Moss**, we do not see in principle why correlations between perceptual and functional properties could not be based on links between the different knowledge stores. It would be good to find a way of putting this to the test.

R2.4. What can we learn from functional brain imaging?

In presenting our argument, we drew not only on neuropsychological cases but also on data from functional brain imaging. **Bub & Bukach**, **Mahon & Caramazza** and **Tyler & Moss** all respond to this. Bub & Bukach and Mahon & Caramazza make a similar point, which is that it is difficult to argue the direction of causation from imaging data. Let us suppose that regions of the brain normally associated with processing the visual properties of objects are active when subjects must make a decision about the biological function of animals. Bub & Bukach and Mahon & Caramazza point out that such evidence is not sufficient for us to argue that these visual properties *need* to become activated to answer these questions – such activations may simply be a by-product of retrieving all forms of knowledge about a stimulus, but they may play no causative role in retrieving functional knowledge. This point is reasonable and well taken. It can be problematic to make causal inferences from imaging data. Nevertheless, the imaging data with normal subjects add to the weight of evidence that should be taken into account when assessing the roles of different types of knowledge in the recognition and naming of particular classes of item. For example, if visual knowledge areas are more activated when answering functional questions about living things than nonliving things (see Thompson-Schill et al. 1999, for this argument), then some form of account of this needs to be given (are all forms of knowledge only retrieved as a by-product for living things?). Also, when functional imaging is combined with neuropsychology, then even stronger arguments can be marshalled. For example, as **Coslett & Saffran** highlight, some imaging studies have been conducted on semantic dementia patients with damage to anterior temporal cortices. Unlike normal subjects, these patients do not activate regions of posterior temporal cortex, associated with high-level visual processing, in associative matching tasks (Mummery et al. 1999), even though the posterior temporal regions were not structurally damaged in the patients. Thus there can be functional changes in neural processing distant to the site of damage. In the relevant studies the investigators did not distinguish effects with different object classes, but if the impact of reduced posterior activation were associated with poor performance with some object classes (e.g., living things), then the argument for a causative link would be strengthened.

Tyler & Moss are concerned with a different point, namely whether there is any evidence from functional imaging studies for regional specialisation in processing different object types, once nuisance variables have been taken account of. We agree with Tyler & Moss that there is a good deal of inconsistency across the reported imaging studies, and we discussed this point in our target article. In an attempt to assess if there are category effects in imaging normal subjects, Devlin et al. (in preparation) have recently re-analysed all the PET data from Price's laboratory in which category was a factor, to evaluate if some common areas do arise across studies. Interesting to note, they do find some consistent areas across studies that vary with object category, with the anterior temporal lobes being more activated for living things and the middle temporal region for tools. This last result is consistent with action knowledge being recruited by tools, since the middle temporal lobe may be part of a circuit involved in action retrieval (Martin et al. 1995; Phillips et al., submitted). Devlin et al. interpret the anterior temporal activation as reflecting the integration of perceptual knowledge within this region, perhaps akin to the notion of a convergence zone (Damasio 1989; see Fig. 4 in our target article). Whatever the valid interpretation, the evidence suggests that there is some neuronal specialisation involved.

R2.5. Underlying mechanisms

One other major theme raised in the commentators is how category effects may emerge. Here it is illuminating to note the suggestions of researchers whose main research interests are not in neuropsychology. **Lamberts**, for example, discusses the contribution that exemplar theories can make in this field, since such models provide comprehensive accounts of data in other areas (e.g., studies of normal categorisation times). He suggests that such theories make the interesting prediction that patients with poor naming of living things may be better able to categorize these items than they can categorize nonliving things. **Lamberts** also notes some data consistent with this approach, though it is not clear that investigators have controlled for factors such as the frequency of the category names or the familiarity of the categories when these comparisons have been made. We look forward to studies that assess categorisation using matched categories.

Tanaka suggests that one reason why category differences may emerge is due to differences in expertise with different object types. He points out that experts tend to list more perceptual features than non-experts, even when the same perceptual features are available to all of us, because experts pay greater attention to the distinctive properties of objects. We accept this point, and we agree that represented "object properties are not only determined by the intrinsic features of an object but also by the experience and motivation of the categorizer." **Tanaka's** point, though, does raise quite a complex set of issues concerning the effects of object familiarity on recognition. Familiarity has

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

typically been considered a nuisance variable that experimenters need to control out. However, as **Lambon Ralph & Garrard** also remark, familiarity can also interact with category, and it may be one of the factors that contributes to gender differences in object identification (Laws 1999). According to **Tanaka's** hypothesis, increased familiarity could lead to the differential weighting of distinctive attributes of objects, making these items less vulnerable to damage. That is, rather than simply being a nuisance variable, familiarity may change the nature of the representations involved. Similarly Lambon Ralph & Garrard suggest that familiar items occupy more semantic space." Clearly the interaction between familiarity and category needs to be explored in more detail rather than controlled away – though, as noted in our target article, familiarity cannot be everything, given that patients can show deficits on objects that they were expert on pre-morbidly (e.g., patient Michelangelo; Sartori & Job 1988).

Several other commentators discuss the role of evolution and learning in determining category-effects. **Barrett**, for example, argues that, from an evolutionary point of view, the perceptual features of living things may be strongly weighted in any representation because perceptual traits are passed on from parents to offspring. This is a slightly different "take" on the evolutionary perspective than that given elsewhere (e.g., Caramazza & Shelton 1998), because it emphasises that weighting of perceptual features may be primary and lead to category effects in adulthood. **Arterberry** also discusses work from developmental psychology which suggests that attention to perceptual properties, such as motion, forms the foundation for category effects that emerge in knowledge representation (see **Zaitchik & Solomon** too). This evolutionary perspective is in line with the HIT framework, which holds that variations in particular kinds of perceptual and functional knowledge, coupled with on-line processing differences, lead to the effects of object category in both patients and normal adults. As Zaitchik & Solomon also stress, we look forward to the convergence between developmental psychology and other strands of cognitive neuroscience in helping us answer questions about how apparent category effects arise (indeed see chapters in Forde & Humphreys, in press).

R2.6. Other categories of knowledge

Our target article dealt with neuropsychological disorders affecting the recognition of living and nonliving things, and not with disorders affecting other categories of knowledge. **Thioux et al.** discuss neuropsychological disorders of number processing, and highlight the contrasting arguments that have been offered to account for such deficits. These arguments parallel some of those found in the literature on differences between living and nonliving things. For instance, Thioux et al. compare a "true categorical account" (numbers are represented independently of other types of knowledge because they are subject to distinct semantic operations) with accounts that stress the fundamental role of other factors – such as visuo-spatial processing or magnitude estimation abilities based on location codes for objects. This mirrors the arguments about whether biological category is crucially important for the living/nonliving distinction, or whether a more basic difference is caused by the contrasts between perceptual and/or functional knowledge in differentiating objects for identification. HIT is aligned with the last position.

Pulvermüller also discusses the role that action-related brain processes may play in the lexical access of words that refer to actions, and he highlights work on impaired lexical access in patients with lesions to motor and adjacent prefrontal areas of the right hemisphere. He suggests that reentrant, action-related information (supported by these regions of the right as well as the left hemisphere) contributes to lexical access for action words. This approach is entirely consistent with our own concerning the processing of living and nonliving things more generally. We hope that some of the predictions that emerge from our account, and some of the issues it raises, can help to provide definitive answers as to why category effects arise in neuropsychology.

ACKNOWLEDGMENT

This work was supported by a grant from the MRC to Glyn Humphreys.

NOTE

1. There may be some exceptions to this, as when a particular stimulus has a particularly diagnostic texture (e.g., a carpet) or a particular pattern of movement.

References

Letters "a" and "r" appearing before authors' initials refer to target article and response, respectively.

Adolph, K. E., Eppler, M. A. & Gibson, E. J. (1993) Development of perception of affordances. In: *Advances in infancy research (Vol. 8),* eds., C. Rovee-Collier & L. P. Lipsitt. Ablex. [MEA]

Adophs, R, Tranel, D., Damasio, H., & Damasio, A. (1994) Impaired recognition of emotion in facial expressions following bilateral damage to the human amygdala. *Nature* 372:669–72. [EMS]

Ahn, W. (1998) Why are different features central for natural kinds and artifacts? The role of causal status in determining feature centrality. *Cognition* 69:135–78. [AP]

Allport, D. A. (1985) Distributed memory, modular systems and dysphasia. In: *Current perspectives in dysphasia,* eds., S. K. Newman & R. Epstein. Edinburgh: Churchill Livingstone. [aGWH]

Arguin, M., Bub, D. & Dudek, G. (1996) Shape integration for visual object recognition and its implication in category specific visual agnosia. *Visual Cognition* 3:221–75. [aGWH, MA, DB]

Arterberry, M. E. & Bornstein, M. H. (in press) Three-month-old infants' categorization of animals and vehicles based on static and dynamic attributes. *Journal of Experimental Child Psychology.* [MEA]

Ashbridge, E. & Perrett, D. I. (1998) Generalizing across object orientation and size. In: *Perceptual constancy. Why things look as they do,* eds. V. Walsh & J. Kulikowski. Cambridge University Press. [MG]

Ashcraft, M.H. (1978) Property dominance and typicality effects in property statement verification. *Journal of Verbal Learning and Verbal Behavior* 17:155–64. [GSC]

Barsalou, L. W. (1985) Ideals, central tendency and frequency of instantiation as determinants of graded structure in categories. *Journal of Experimental Psychology: Learning, Memory and Cognition* 11:629–54. [ECL]

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

Behl-Chadha, G. (1996) Basic-level and superordinate-like categorical representations in early infancy. *Cognition* 60:105–41. [MEA]

Berlin, B. (1999) How a folkbotanical system can be both natural and comprehensive: One Maya Indian's view of the plant world. In: *Folkbiology,* eds., D. L. Medin & S. Atran. MIT Press. [JD]

Biederman, I. (1987) Recognition-by-components: A theory of human image understanding. *Psychological Review* 94:115–47. [rGWH, MA, MG]

Biederman, I. & Cooper, E. E. (1991) Priming contour-deleted images: Evidence

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

for intermediate representations in visual object recognition. *Cognitive Psychology* 23:393–419. [MG]
Biederman, I. & Gerhardstein, P. C. (1993) Recognizing depth-rotated objects: Evidence and conditions for three-dimensional viewpoint invariance. *Journal of Experimental Psychology: Human Perception and Performance* 19:1162–82. [MG]
Biederman, I. & Ju, G. (1988) Surface versus edge-based determinants of visual recognition. *Cognitive Psychology* 20:38–64. [JD]
Bisiach, E. (1966) Perceptual factors in the pathogenesis of anomia. *Cortex* 2:90–5. [JD]
Bloom, P. (1996) Intention, history, and artifact concepts. *Cognition* 60:1–29. [HCB]
Bonda, E., Petrides, M., Ostry, D. & Evans, A. (1996) Specific involvement of human parietal systems and the amygdala in the perception of biological motion. *Journal of Neuroscience* 16:3737–44. [aGWH]
Bornstein, M. H. (1984) A descriptive taxonomy of psychological categories used by infants. In: *Origins of cognitive skills,* ed., C. Sophian. Lawrence Erlbaum. [MEA]
Boster, J. S. & Johnson, J. C. (1989) Form or function: A comparison of expert and novice judgments of similarity among fish. *American Anthropologist* 91:866–89. [JT]
Boucart, M. & Humphreys, G. W. (1992) The computation of perceptual structure from collinearity and closure: Normality and pathology. *Neuropsychologia* 30:527–46. [aGWH]
Bozeat, S., Lambon Ralph, M. A., Patterson, K., Garrard, P. & Hodges, J. R. (2000) Non-verbal semantic impairment in semantic dementia. *Neuropsychologia* 38:1207–15. [MALR]
Braitenberg, V. & Pulvermueller, F. (1992) Entwurf einer neurologischen Theorie der Sprache. *Naturwissenschaften* 79:103–17. [FP]
Breedin, S. D., Saffran, E. M. & Coslett, H. B. (1994) Reversal of the concreteness effect in a patient with semantic dementia. *Cognitive Neuropsychology* 11:617–60. [aGWH, EMS]
Bub, D. (2000) Methodological issues confronting PET and fMRI studies of cognitive function. *Cognitive Neuropsychology* 17:467–84. [DB]
Burton, A. M. (1994) Learning new faces in an interactive activation and competition model. *Visual Cognition* 1:313–48. [aGWH]
Buxbaum, L. & Coslett, H. B. (in press) Specialized structural descriptions for human body parts: Evidence from autotopagnosia. *Neuropsychologia.* [EMS]
Campbell, R. & Manning, L. (1996) Optic aphasia: A case with spared action naming and associated disorders. *Brain and Language* 53:183–221. [EMS]
Cappa, S. F., Frugoni, M., Pasquali, P., Perani, D. & Zorat, F. (1998) Category-specific naming impairment for artefacts: A new case. *Neurocase* 4:391–97. [aGWH]
Caramazza, A. (1998) The interpretation of semantic category-specific deficits: What do they reveal about the organisation of conceptual knowledge in the brain? *Neurocase* 4:265–72. [aGWH, BM]
(2000) Minding the facts: A comment on Thompson-Schill et al.'s "A neural basis for category and modality specificity of semantic knowledge." *Neuropsychologia* 38:944–49. [BM]
Caramazza, A. & Shelton, J. R. (1998) Domain specific knowledge systems in the brain: The animate-inanimate distinction. *Journal of Cognitive Neuroscience* 10:1–34. [arGWH, ECL, BM]
Caramazza, A., Hillis, A. E., Rapp, B. C. & Romani, C. (1990) The multiple semantics hypothesis: Multiple confusions? *Cognitive Neuropsychology* 7:161–89. [aGWH, EMS]
Carey, S. (1985) *Conceptual change in childhood.* Bradford Books. [DZ]
Carnap, R. (1956). *Meaning and necessity,* 2nd ed. Chicago University Press. [RGdeA]
Cave, C. B. & Kosslyn, S. M. (1993) The role of parts and spatial relations in object identification. *Perception* 22:229–48. [MG]
Chan, A., Butters, N. & Salmon, D. (1997) The deterioration of semantic networks in patients with Alzheimer's disease: A cross sectional study. *Neuropsychologica* 35:241–48. [DZ]
Chao, L. L., Haxby, J. V. & Martin, A. (1999) Attribute-based neural substrate in posterior temporal cortex for perceiving and knowing about objects. *Nature Neuroscience* 2:913–19. [aGWH, BM, LKT]
Chertkow, H., Bub, D. & Caplan, D. (1992) Constraining theories of semantic memory processing: Evidence from dementia. *Cognitive Neuropsychology* 9:273–368. [aGWH]
Cipolotti, L., Butterworth, B. & Denes, G. (1991) A specific deficit for numbers in a case of dense acalculia. *Brain* 114:2619–37. [MT]
Cohen, G. & Burke, D. M. (1993) *Memory for proper names.* Lawrence Erlbaum. [aGWH]
Coley, J. D., Medin, D. L. & Atran, S. (1997) Does rank have its privilege? Inductive inferences within folkbiological taxonomies. *Cognition* 64:73–112. [HCB]
Cotheart M., Inglis L., Cupples L., Michie P., Bates A. & Budd B. (1998) A semantic subsystem specific to the storage of information about the visual attributes of animate and inanimate objects. *Neurocase* 4:353–70. [BM]
D'Esposito, M., Detre, J. A., Aguire, G. K., Aslop, D. C., Tippett, L. J. & Farah, M. J. (1997) A functional MRI study of mental image generation. *Neuropsychologia* 35:725–30. [aGWH]
Damasio, A. R. (1989) Time-locked multiregional retroactivation: A systems-level proposal for the neural substrates of recall and recognition. *Cognition* 33:25–62. [arGWH]
(1990) Category-related recognition defects as a clue to the neural substrates of knowledge. *Trends in Neuroscience* 13:95–98. [aGWH]
Damasio, H., Grabowski, T. J., Tranel, D., Hitchwa, R. D. & Damasio, A. R. (1996) A neural basis for lexical retrieval. *Nature* 380:499–505. [aGWH, LKT]
Davidoff, J. (1991) *Cognition through Color.* MIT Press. [JD]
Davidoff, J. & de Bleser, R. (1993) Optic aphasia: A review of past studies and reappraisal. *Aphasiology* 7:135–54. [JD]
Davidoff, J. & Ostergaard, A. L. (1988) The role of colour in categorical judgements. *Quarterly Journal of Experimental Psychology* 40A:533–44. [JD]
Davidoff, J. & Warrington. E. K. (1999) The bare bones of object recognition: implications from a case of object recognition impairment. *Neuropsychologia* 37:279–92. [rGWH, JD]
De Almeida, R. G. (1999a) What do category-specific semantic deficits tell us about the representation of lexical concepts? *Brain and Language* 68:241–48. [RGdeA]
(1999b) The representation of lexical concepts: A psycholinguistic inquiry. Unpublished Ph.D. dissertation. Rutgers University, New Brunswick. [RGdeA]
De Almeida, R. G. & Fodor, J. A. (1996) Still looking for structural complexity effects in the representation of lexical concepts. In: *Proceedings of the Eighteenth Annual Conference of the Cognitive Science Society,* ed. G. Cotrell. Lawrence Erlbaum. [RGdeA]
De Renzi, E. & Lucchelli, F. (1994) Are semantic systems separately represented in the brain? The case of living category impairment. *Cortex* 30:3–25. [aGWH]
Decety, J., Perani, D., Jeannerod, M., Bettinardi, V., Tadary, B., Woods, R., Mazziotta, J. C. & Fazio, F. (1994) Mapping motor representations with positron emission tomography. *Nature* 371:600–602. [aGWH]
Dehaene, S. (1995) Electrophysiological evidence for category-specific word processing in the normal human brain. *NeuroReport* 6:2153–57. [aGWH]
(1997) *The number sense.* Oxford University Press. [MT]
Dehaene, S. & Changeux, J-P. (1993) Development of elementary numerical abilities: A neuronal model. *Journal of Cognitive Neuroscience* 5:390–407. [MT]
Dehaene, S., Spelke, E., Pinel, P., Stanescu, R. & Tsivkin, S. (1999) Sources of mathematical thinking: Behavioral and brain-imaging evidence. *Science* 284:970–74. [MT]
Dennis, M. (1976) Dissociated naming and locating of body parts after left anterior temporal lobe resection: An experimental case study. *Brain and Language* 3:147–63.
Devlin, J. T., Gonnerman, L., Anderson, E. & Seidenberg, M. (1998) Category-specific deficits in focal and widespread damage: A computational account. *Journal of Cognitive Neuroscience* 10:77–94. [arGWH, ECL]
Devlin, J. T., Moore, C. J., Mummery, C. J., Gorno-Tempini, M., Phillips, J. A., Nopenney, U. & Price, C. J. (in preparation). Anatomic constraints on cognitive theories of category specificity. [rGWH]
Devlin, J. T., Russell, R., Davis, M., Price, C., Moss, H., Fadili, J. & Tyler, L. K. (in press) Is there an anatomical basis for category-specificity? Semantic memory studies in PET and fMRI. *Neuropsychologia.* [LKT]
Disendruck , G. & Gelman, S. A. (1999) Domain differences in absolute judgments of category membership: Evidence for an essentialist account of categorization. *Psychonomic Bulletin and Review* 6:338–46. [JD]
Dixon, M., Bub, D. N., & 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* 14:1085–130. [aGWH, MA, DB]
(1998). Semantic and visual determinants of face recognition in a prosopagnosic patient. *Journal of Cognitive Neuroscience* 10:362–76. [aGWH]
Durrant-Peatfield, M., Tyler, L., Moss, H., & Levy, J. (1997) The distinctiveness of form and function in category structure: A connectionist model. In: *Proceedings of the 19th Annual Conference of the Cognitive Science Society.* Lawrence Erlbaum. [arGWH, LKT]
Farah, M. J. (2000) *The cognitive neuroscience of vision.* Blackwell. [MG]
Farah , M. J. & McClelland, J. L. (1991) A computational model of semantic memory impairment: Modality specificity and emergent category specificity. *Psychological Review* 120:339–57. [aGWH, MA, HCB, AP]
Farah, M. J. & Wallace, M. A. (1992) Semantically-bounded anomia: Implications for the neural implementation of naming. *Neuropsychologia* 30:609–21. [aGWH, BM]

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

Farah, M. J., Meyer, M. M. & McMullen, P. A. (1996) The living/nonliving dissociation is not an artifact: Giving an a priori implausible hypothesis a strong test. *Cognitive Neuropsychology* 13:137–54. [aGWH]

Fodor, J. A. (1990) *A theory of content and other essays.* MIT Press. [RGdeA]

(1998) *Concepts: Where cognitive science went wrong.* Oxford University Press. [RGdeA]

Fodor, J. D., Fodor, J. A. & Garrett, M. F. (1975) The psychological unreality of semantic representations. *Linguistic Inquiry* 6:515–31. [RGdeA]

Forde, E. M. E. (1999) Category specific recognition impairments. In: *Case studies in the neuropsychology of vision,* ed., G. W. Humphreys. Psychology Press. [aGWH]

Forde, E. M. E., Francis, D., Riddoch, M. J., Rumiati R. & Humphreys, G. W. (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(3):403–58. [arGWH, JD, KL]

Forde, E. M. E. & Humphreys, G. W. (1995) Refractory semantics in global aphasia: On semantic organization and the access-storage distinction in neuropsychology. *Memory* 3:265–307. [aGWH]

(1997) A semantic locus for refractory behaviour: Implications for access-storage distinctions and the nature of semantic memory. *Cognitive Neuropsychology* 14:367–402. [aGWH]

(1999) Category-specific recognition impairments: A review of important case studies and influential theories. *Aphasiology* 13:169–93. [aGWH]

(in press) *Category specificity in brain and mind.* Psychology Press. [rGWH]

Funnell, E. & De Mornay Davies, P. (1996) JBR: A reassessment of concept familiarity and a category-specific disorder for living things. *Neurocase* 2:461–74. [aGWH, MALR]

Funnell, E. & Sheridan, J. (1992) Categories of knowledge? Unfamiliar aspects of living and nonliving things. *Cognitive Neuropsychology* 9:135–53. [aGWH, LKT]

Gaffan, D. & Heywood, C. A. (1993) A spurious category-specific visual agnosia for living things in normal human and nonhuman primates. *Journal of Cognitive Neuroscience* 5:118–28. [arGWH, KRL]

Gainotti, G. & Silveri, M. C. (1996) Cognitive and anatomical locus of lesion in a patient with a category-specific semantic impairment for living beings. *Cognitive Neuropsychology* 13:357–90. [MALR]

Gainotti, G., Silveri, M. C., Daniele, A. & Giustolisi, L. (1995) Neuroanatomical correlates of category-specific semantic disorders: A critical survey. *Memory* 3:247–64. [aGWH, MALR]

Gale, T. M., Done, D. J. & Frank, R. J. (in press) Visual crowding and category specific deficits: A neural network model. *Cognitive Neuropsychology.* [rGWH]

Garrard, P., Lambon Ralph, M. A., Hodges., J. R. & Patterson, K. (in press) Prototypicality, distinctiveness, and intercorrelation: Analyses of the semantic attributes of living and nonliving concepts. *Cognitive Neuropsychology.* [BM]

Garrard, P., Lambon Ralph, M. A. & Hodges, J. R. (in press a) Semantic dementia: A category-specific paradox. In: *Category-specificity in brain and mind,* eds., E. M. E. Forde & G. W. Humphreys. Psychology Press. [MALR]

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* 18:125–74 [MALR]

Garrard, P., Lambon Ralph, M. A., Watson, P., Powis, J., Patterson, K. & Hodges, J. R. (in press b) Longitudinal profiles of semantic impairment for living and nonliving concepts in dementia of Alzheimer's type. *Journal of Cognitive Neuroscience.* [MALR]

Garrard, P., Patterson, K., Watson, P. C. & Hodges, J. R. (1998) Category-specific semantic loss in dementia of Alzheimer's type. *Brain* 121:633–46. [aGWH, MALR]

Gelman, R. (1990) First principles organize attention to and learning about relevant data: Number and the animate-inanimate distinction as examples. *Cognitive Science* 14:79–106. [aGWH, MEA]

Gelman, S. (1988) The development of induction within natural kind and artifact categories. *Cognitive Psychology* 20:65–95. [aGWH]

Gelman, S. & Kremer, K. E. (1991) Understanding natural cause: Children's explanations of how objects and their properties originate. *Child Development* 62:396–414. [aGWH]

Genkinger, B. & Pulvermüller, F. (in press) The right hemisphere's role in action verb processing. *Neurocase.* [FP]

Gerlach, C. (in press) Structural similarity causes different category-effects depending on task characteristics. *Neuropsychologia.* [rGWH]

Gibson, J. J. (1979) *The ecological approach to visual perception.* Houghton Mifflin. [rGWH, MEA]

Glaser, W. R. (1994) Picture naming. *Cognition* 42:61–105. [aGWH]

Goldstone, R. L. (1996) Alignment-based nonmonotonicities in similarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 22:988–1001. [MG]

Gonnerman, L. M., Andersen, E. S., Devlin, J. T., Kempler, D. & Seidenberg, M. S. (1997) Double dissociation of semantic categories in Alzheimer's disease. *Brain and Language* 57:254–79. [aGWH, MALR]

Goodman, N. (1972). Seven strictures on similarity. In: *Problems and projects,* ed., N. Goodman. Bobbs-Merrill. [ECL]

Grabowski, T. J., Damasio, H. & Damasio, A. R. (1998) Premotor and prefrontal correlates of category-related lexical retrieval. *Neurimage* 7:232–43. [aGWH]

Grafton, S. T., Fadiga, L., Arbib, M. A. & Rizzolatti, G. (1997) Premotor cortex during observation and naming of familiar tools. *Neuroimage* 6:231–36. [aGWH]

Guistolisi, L., Bartolomeno, P., Daniele, A., Marra, C. & Gainotti, G. (1993) Category-specific semantic impairment for living things in the early stages of Alzheimer's disease: Further evidence from a study on single cases. *Journal of Clinical and Experimental Neuropsychology* 15:403.

Hart, J. & Gordon, B. (1992) Neural subsystems for object knowledge. *Nature* 359:60–64. [aGWH, ECL, BM]

Hart, J., Berndt, R. S. & Caramazza, A. (1985) Category-specific naming deficit following cerebral infarction. *Nature* 316:439–40. [aGWH, ECL, BM]

Hatano, G. & Inagaki, K. (1994) Young children's naïve theory of biology. *Cognition* 50:171–88. [DZ]

Hebb, D. O. (1949) *The organization of behavior.* Wiley.

Hillis, A. E. & Caramazza, A. (1991) Category-specific naming and comprehension impairment: A double dissociation. *Brain* 114:2081–94. [aGWH, ECL, BM]

(1995) Constraining claims about theories of semantic memory: More on unitary versus multiple semantics. *Cognitive Neuropsychology* 12:175–86. [aGWH]

Howard, D. & Patterson, K. (1992) *The pyramids and elm trees test.* Thames Valley Test Company. [EMS]

Hummel, J. E. & Biederman, I. (1992) Dynamic binding in a neural network for shape recognition. *Psychological Review* 99:480–517. [MG]

Humphrey, G. K. & Jolicoeur, P. (1993) An examination of the effects of axis foreshortening, monocular depth cues, and visual field on object identification. *The Quarterly Journal of Experimental Psychology* 46A:137–59. [MG]

Humphreys, G. W. & Price, C. J. (in press) Cognitive neuropsychology and functional brain imaging: Implications for functional and anatomical models of cognition. *Psychological Research.* [rGWH]

Humphreys, G. W. & Riddoch, M. J. (1999) Impaired development of semantic memory: Separating semantic from structural knowledge and diagnosing a role for action in establishing stored memories for objects. *Neurocase* 5:519–32. [aGWH]

Humphreys, G. W. & Rumiati, R. (1998) When joys come not in single spies but in battalions: Within category and within modality identification increases the accessibility of degraded semantic knowledge. *Neurocase* 4:111–26. [aGWH]

Humphreys, G. W., Francis, D. & Samson, D. (in preparation) Category specific naming deficit linked to impaired visual knowledge. [aGWH]

Humphreys, G. W., Lamote, C. & Lloyd-Jones, T. J. (1995) An interactive activation approach to object processing: Effects of structural similarity, name frequency, and task in normality and pathology. *Memory* 3:535–86. [arGWH, BM]

Humphreys, G. W., Price, C. & Riddoch, M. J. (1999) From objects to names: A cognitive neuroscience approach. *Psychological Research* 62:118–30. [arGWH]

Humphreys, G. W., Riddoch, M. J. & Price, C. (1997) Top-down processes in object identification: Evidence from experimental psychology, neuropsychology and functional anatomy. *Proceedings for the Royal Society B* 352:1275–82. [aGWH]

Humphreys, G. W., Riddoch, M. J. & Quinlan, P. T. (1985) Interactive processes in perceptual organization: Evidence from visual agnosia. In: *Attention and performance XI,* eds., M. I. Posner & O. S. M. Marin. Lawrence Erlbaum. [aGWH]

(1988) Cascade processes in picture identification. *Cognitive Neuropsychology* 5:67–103. [arGWH, MA, JD, MG, LKT]

Jackendoff, R. (1990) *Semantic structures.* MIT Press. [ECL]

Job, R., Rumiati, R. & Lotto, L. (1992) The picture superiority effect in categorization: Visual or semantic? *Journal of Experimental Psychology: Learning, Memory and Cognition* 18:1019–28. [aGWH]

Jolicoeur, P. & Humphrey, G. K. (1998) Perception of rotated two-dimensional and three-dimensional objects and visual shapes. In: *Perceptual constancy. Why things look as they do,* eds., V. Walsh & J. Kulikowski. Cambridge University Press. [MG]

Keil, F. (1987) Conceptual development and category structure. In: *Concepts and conceptual development: Ecological and intellectual factors in categorisation,* ed., U. Neisser. Cambridge University Press. [aGWH]

Kiefer, M. (in press) Perceptual and semantic sources of category-specific effects: Event-related potentials during picture and word categorization. *Perception and Psychophysics.* [aGWH]

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

Kohonen, T. (1988) *Self-organisation and associative memory.* Springer-Verlag. [rGWH]

Kruschke, J. K. (1992) ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review* 99:22–44. [aGWH, KL]

Kurbat, M. A. (1997) Can the recognition of living things really be selectively impaired? *Neuropsychologia* 35:813–27. [aGWH]

Kurbat, M. A. & Farah, M. J. (1998) Is the category-specific deficit for living things spurious? *Journal of Cognitive Neuroscience* 10:355–61. [aGWH]

Laiacona, M., Barbarotto, R. & Capitani, E. (1993) Perceptual and associative knowledge in category specific impairment of semantic memory: A study of two cases. *Cortex* 29:727–40. [aGWH, BM]

Laiacona, M., Capitani, E. & Barbarotto, R. (1997) Semantic category dissociations: A longitudinal study of two cases. *Cortex* 33:441–61. [aGWH]

Lamberts, K. (1995) Categorization under time pressure. *Journal of Experimental Psychology: General* 124:161–80. [KL]

(2000) Information-accumulation theory of speeded categorization. *Psychological Review* 107:227–60. [KL]

Lamberts, K. & Shapiro, L. (in press) Exemplar models and category-specific deficits. In: *Category-specificity in brain and mind,* eds., E. M. E. Forde and G. W. Humphreys. Psychology Press. [KL, rGWH]

Lambon Ralph, M. A., Graham, K. S., Ellis, A. W. & Hodges, J. R. (1998a) Naming in semantic dementia–What matters? *Neuropsychologia* 36:775–84. [MALR]

Lambon Ralph, M. A., Graham, K. S., Patterson, K. & Hodges, J. R. (1999) Is a picture worth a thousand words? Evidence from concept definitions by patients with semantic dementia. *Brain and Language* 70:309–35. [MALR]

Lambon Ralph, M. & Howard, D. (2000). Gogi aphasia or semantic dementia? Simulating and accessing poor verbal comprehension in a case of progressive verbal aphasia. *Cognitive Neuropsychology* 17:437–66.

Lambon Ralph, M. A., Howard, D., Nightingale, G. & Ellis A. W. (1998b) Are living and non-living category-specific deficits causally linked to impaired perceptual or associative knowledge? Evidence from a category-specific double dissociation. *Neurocase* 4:311–38. [arGHW, BM]

Lambon Ralph, M. A., McClelland, J., Patterson, K., Galton, C. & Hodges, J. (in press) No right to speak? The relationship between object naming and semantic impairment: Neuropsychological evidence and a computational model. *Journal of Cognitive Neuroscience.* [MALR]

Lambon Ralph, M. A., Patterson, K. & Hodges, J. R. (1997) The relationship between naming and semantic knowledge for different categories in dementia of Alzheimer's type. *Neuropsychologia* 35:1251–60. [MALR]

Lauro-Grotto, R., Piccini, C. & Shallice, T. (1997) Modality-specific operations in semantic dementia. *Cortex* 33:593–622. [EMS]

Laws, K. R. (1999) Gender affects latencies for naming living and nonliving things. *Cortex* 35:729–33. [KRL]

(2000). Category-specific naming errors in normal subjects: the influence of evolution and experience. *Brain and Language* 75:123–33. [rGWH, KRL]

Laws, K. R., Evans, J. J., Hodges, J. R. & McCarthy, R. A. (1995) Naming without knowing and appearance without associations: Evidence for constructive processes in semantic memory? *Memory* 3:409–33. [JD]

Laws, K. R. & Gale, T. M. (in press) Category-specific naming and the "visual" characteristics of line drawn stimuli. *Cortex.* [KRL]

Laws, K. R. & Gale, T. M., Frank, R. & Davey, N. (in press) Visual similarity is greater for line drawings of nonliving than living things: The importance of musical instruments and body-parts. *Brain and Cognition.* [KRL]

Laws, K. R. & Neve, C. (1999) A "normal" category-specific advantage for naming living things. *Neuropsychologia* 37:1263–69. [rGWH, MEA, KRL]

Lawson, R. (1999) Achieving visual object constancy across plane rotation and depth rotation. *Acta Psychologica 102*:221–45. [MG]

Lawson, R. & Humphreys, G. W. (1998) View-specific effects of depth-rotation and foreshortening on the initial recognition and priming of familiar objects. *Perception and Psychophysics* 60:1052–66. [MG]

Lawson, R., Humphreys, G. W. & Watson, D. G. (1994) Object recognition under sequential viewing conditions: Evidence for viewpoint-specific recognition procedures. *Perception* 23:595–614. [MG]

Leek, E. C. (in press) Category-specific semantic memory impairments: What can connectionist simulations reveal about the organisation of conceptual knowledge? In: *Connectionist models in cognitive psychology,* ed., G. Houghton. Psychology Press. [ECL]

Lewandowsky, M. (1908) Ueber Absphaltung des Farbensinnes. *Monatsschrift fuer Psychiatrie und Neurologie* 23:488–510 (see Davidoff, J. & Fodor, G. [1989] An annotated translation of Lewandowsky [1908] in *Cognitive Neuropsychology* 6:165–77). [JD]

Lissauer, H. (1890) Ein Fall von Seelenblindheit nebst einem Beitrag zur Theorie derselben. *Archiv fur Psychiatrie* 21:222–270. Edited and reprinted in translation by Jackson, M. (1988) Lissauer on agnosia. *Cognitive Neuropsychology* 5:157–92. [aGWH]

Lloyd- Jones, T. & Humphreys, G. W. (1997) Perceptual differentiation as a source of contour effects in object processing: Evidence from naming and object decision. *Memory and Cognition* 25:18–35. [aGWH, KRL]

Luzzatti, C. & Davidoff, J. (1994) Impaired retrieval of object-colour knowledge with preserved our naming. *Neuropsychologia* 32:1–18. [arGWH, BM, JD]

Madole, K. & Cohen, L. (1995) The role of object parts in infants' attention to form-function correlations. *Developmental Psychology* 31:637–48. [MEA]

Malt, B. C. & Johnson, E. C. (1992) Do artifact concepts have cores? *Journal of Memory and Language* 31:195–217. [AP]

Mandler, J. M. (1992) How to build a baby: 2. Conceptual primitives. *Psychological Review* 99:587–604. [aGWH]

(2000) Perceptual and conceptual processes in infancy. *Journal of Cognition and Development* 1:3–36. [MEA]

Mandler, J. M., Bauer, P. J. & McDonough, L. (1991) Separating the sheep from the goats: Differentiating global categories. *Cognitive Psychology* 23:263–98. [aGWH]

Mandler, J. M. & McDonough, L. (1996) Drinking and driving don't mix: Inductive generalisation in infancy. *Cognition* 59:307–35. [aGWH]

Marr, D. (1982) Vision: A computational investigation into the human representation and processing of visual information. W. H. Freeman. [aGWH]

Marr, D. & Nishihara, H. K. (1978) Representation and recognition of the spatial organization of three-dimensional shapes. *Proceedings of the Royal Society, London B* 200:269–94. [MG]

Marshall, J., Pring, T., Chiat, S. & Robson, J. (1996) Calling a salad a federation: An investigation of semantic jargon. Part I–Nouns. *Journal of Neurolinguistics* 9:237–50. [EMS]

Martin, A., Haxby, J. V., Lalonde, F. M., Wiggs, C .L. & Ungerleider, L. G. (1995) Discrete cortical regions associated with knowledge of color and knowledge of action. *Science* 270:102–105. [arGWH]

Martin, A., Wiggs, C. L., Ungerleider, L. G., & Haxby, J. V. (1996) Neural correlates of category-specific knowledge. *Nature* 379:649–52. [aGWH, LKT]

Massey, C. & Gelman, R. (1988) Preschoolers' ability to decide whether pictured unfamiliar objects can move themselves. *Developmental Psychology* 24:307–17. [aGWH]

McCarthy, R. (1995) *Semantic knowledge and semantic representations.* Taylor and Francis. [aGWH]

McCarthy, R. & Warrington, E. K. (1994) Disorders of semantic memory. *Philosophical Transactions of the Royal Society of London* 346:89–96. [aGWH]

McClelland, J. L. & Rumelhart, D. E. (1985) Distributed memory and the representation of general and specific information. *Journal of Experimental Psychology, General* 114:159–88. [aGWH]

McKenna, P. & Warrington, E. K. (1978) Category-specific naming preservation: A single case study. *Journal of Neurology, Neurosurgery, and Psychiatry* 41:571–74. [ECL]

McRae, K. & Cree, G. S. (2000) Factors underlying category-specific semantic deficits. In: *Category-specificity in mind and brain,* eds., E. M.E. Forde & G. W. Humphreys. Psychology Press. [GSC, BM, MALR]

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* 53:360–73. [GSC]

McRae, K., de Sa, V. R. & Seidenberg, M. S. (1997) On the nature and scope of featural representations of word meaning. *Journal of Experimental Psychology, General* 126:99–130. [aGWH, MA, ECL]

Medin, D. L. & Atran, S. (1999) *Folkbiology.* MIT Press. [DZ]

Medin, D. L., Lynch, E. B., Coley, J. D. & Atran, S. (1997) Categorization and reasoning among tree experts: Do all roads lead to Rome? *Cognitive Psychology* 32:49–96. [JT]

Miceli, G., Fouch, E., Capasso, R., Shelton, J. R., Tomaiuolo, F. & Caramazza, A. (in press) The dissociation of color from form and function knowledge. *Nature Neuroscience.* [BM]

Mikkulainen, R. (1997) Dyslexic and category-specific aphasic impairments in a self-organising feature map model of the lexicon. *Brain and Language* 59:334–66. [aGWH]

Miller, G. A. (1981) *Language and speech.* Freeman. [MEA]

Montanes, P., Goldblum, M. C. & Boller, F. (1995) The naming impairment of living and nonliving items in Alzheimer's disease. *Journal of the International Neuropsychological Society* 1:39–48. [JD]

Moore, C. J. & Price, C. J. (1999) A functional neuroimaging study of the variables that generate category-specific object processing differences. *Brain* 122:943–62. [aGWH]

Morton, J. & Patterson, K. (1980) A new attempt at an interpretation, or, an attempt at a new interpretation. In: *Deep dyslexia,* eds., M. Coltheart, K. Patterson and J. C. Marshall. Routledge and Kegan Paul. [aGWH]

Moss, H. E. & Tyler, L. K. (2000) A progressive category-specific semantic deficit for non-living things. *Neuropsychologia* 38:60–82. [LKT]

Moss, H. E., Tyler, L. K. & Devlin, J. T. (in press) Emerging categories:

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

Conceptual structure in a distributed system. *Journal of Memory and Language.* [aGWH]
Moss, H. E., Tyler, L. K. & Jennings, F. (1997) When leopards lose their spots: Knowledge of visual properties in category-specific deficits for living things. *Cognitive Neuropsychology* 14:901–50. [aGWH, MA, ECL]
Moss, H. E., Tyler, L. K., Durrant-Peatfield, M. & Bunn, E. M. (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. [aGWH, BM, KL]
Mummery, C. J., Patterson, K., Hodges, J. R. & Wise, R. (1996) Generating "tiger" as an animal name or a word beginning with T: Differences in brain activation. *Proceedings of the Royal Society of London Series B* 263:989–95. [LKT]
Mummery, C. J., Patterson, K., Price, C. J., Ashburner, J., Frackowiak, R. S. J. & Hodges, J. R. (2000) A voxel based morphometry study of semantic dementia: The relation of temporal lobe atrophy to cognitive deficit. *Annals of Neurology* 47:36–45. [MALR]
Mummery, C. J., Patterson, K., Wise, R. J. S., Vandenbergh, R.. Price, C. J. & Hodges, J. R. (1999) Disrupted temporal lobe connections in semantic dementia. *Brain* 122:61–73. [rGWH]
Murphy, G. L. (1991) Parts in object concepts: Experiments with artificial categories. *Memory and Cognition* 19:423–38. [MG]
Nielsen, J. M. (1946) *Agnosia, apraxia, aphasia: Their value in cerebral localization,* 2nd edition. Hoeber. [aGWH]
Nosofsky, R. M. (1986) Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General* 115:39–57. [KL]
(1987) Attention and learning processes in the identification and categorization of integral stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 13:87–109. [KL]
Nosofsky, R. M. & Palmeri, T. J. (1997) An exemplar-based random walk model of speeded classification. *Psychological Review* 104:266–300. [KL]
Nosofsky, R. M. & Zaki, S. R. (1998) Dissociations between categorization and recognition in amnesic and normal individuals: An exemplar-based interpretation. *Psychological Science* 9:247–55. [KL]
Ostergaard, A. L. & Davidoff, J. (1988) Some effects of color on naming and recognition of objects. *Journal of Experimental Psychology: Learning, Memory and Cognition* 11:579–87. [JD]
Paivio, A. (1991) Dual coding theory: Retrospect and current status. *Canadian Journal of Psychology* 45:255–87. [EMS]
Palmer, S. E. (1977) Hierarchical structure in perceptual representation. *Cognitive Psychology* 9:441–74. [MG]
Perani, D., Cappa, S., Bettinardi, V., Bressi, S., Gornotempini, M., Matarrese, M. & Fazio, F. (1995) Different neural systems for the recognition of animals and man-made tools. *NeuroReport* 6:1637–41. [aGWH, LKT]
Perry, C. (1999) Testing a computational account of category-specific deficits. *Journal of Cognitive Neuroscience* 11:312–20. [aGWH]
Phillips, J. A., Humphreys, G. W., Nopenney, U. & Price, C. J. (submitted) The neural substrates of action retrieval: An examination of semantic and visual routes to action. *Visual Cognition.*
Pinel, P., Le Clerc H, G., van de Moortele, P-F., Naccache, L., Le Bihan, D. & Dehaene, S. (1999) Event-related fMRI analysis of the cerebral circuit for number comparison. *Neuroreport* 10:1473–79. [MT]
Potter, M. C. & Faulconer, B. A. (1975) Time to understand pictures and words. *Nature* 253:437–38. [aGWH]
Poulin-Dobois, D., Lepage, A. & Ferland, D. (1996) Infants' concept of animacy. *Cognitive Development* 11:19–36. [aGWH]
Powell, J. & Davidoff, J. (1995) Selective impairments of object knowledge in a case of acquired cortical blindness. *Memory* 3:435–61. [JD]
Price, C. J. & Humphreys, G. W. (1989) The effects of surface detail on object categorization and naming. *Quarterly Journal of Experimental Psychology* 41A:797–827. [arGWH, JD]
Price, C. J., Moore, C. J., Humphreys, G. W., Frackowiak, R. S. J. & Friston, K. J. (1996) The neural regions sustaining object recognition and naming. *Proceedings of the Royal Society of London, Series B* 263:1501–1507. [aGWH, DB, LKT]
Pulvermüller, F. (1992) Constituents of a neurological theory of language. *Concepts in Neuroscience* 3:157–200. [FP]
(1996). Hebb's concept of cell assemblies and the psychophysiology of word processing. *Psychophysiology* 33:317–33. [FP]
(1999). Words in the brain's language. *Behavioral and Brain Sciences* 22:253–336. [aGWH, FP]
Pulvermüller, F. Ehlert, C. & Elbert, T. (1998) Processing of action words is impaired in right-handers with left-sided hemiparesis. *Society of Neuroscience Abstracts* 22:15. [FP]
Pulvermüller, F., Haerle, M. & Hummel, F. (2000) Neurophysiological distinction of verb categories. *&Neuroreport* 11:2789–93. [FP]
Pulvermüller, F., Lutzenberger, W. & Preissl, H. (1999) Nouns and verbs in the intact brain: Evidence from event-related potentials and high-frequency cortical responses. *Cerebral Cortex* 9:498–508. [FP]
Pulvermüller, F. & Preissl, H. (1991) A cell assembly model of language. *Network: Computation in Neural Systems* 2:455–68. [FP]
Pulvermüller, F., Preissl, H., Lutzenberger, W. & Birbaumer, N. (1996) Brain rhythms of language: Nouns vs. verbs. *European Journal of Neuroscience* 8:937–41. [aGWH]
Quinn, P. C. & Eimas, P. D. (1996) Perceptual cues that permit categorical differentiation of animal species by infants. *Journal of Experimental Child Psychology* 63:189–211. [MEA]
Rakison, D. H. & Butterworth, G. E. (1998) Infants' use of object parts in early categorization. *Developmental Psychology* 34:49–62. [MEA]
Rakison, D. H. & Poulin-Dubois, D. (in press) The developmental origin of the animate-inanimate distinction. *Psychological Bulletin.*
Rapcsak, S. Z., Kaszniak, A. W. & Rubens, A. B. (1989) Anomia for facial expressions: Evidence for a category specific visual-verbal disconnection syndrome. *Neuropsychologia* 27:1031–41. [ECL]
Rapp, B. C. & Caramazza, A. (1991) Lexical deficits. In: *Acquired aphasias, 2nd ed.*, M. T. Sarno. Academic Press. [RGdeA]
Riddoch, M. J. & Humphreys, G. W. (1987a) Picture naming. In: *Visual object processing: A cognitive neuropsychological approach,* eds., G. W. Humphreys & M. J. Riddoch. LEA. [aGWH]
(1987b) A case of integrative visual agnosia. *Brain* 110:1431–62. [aGWH]
(1987c) Visual object processing in optic aphasia: A case of semantic access agnosia. *Cognitive Neuropsychology* 4:131–85. [aGWH]
(1993) The smiling giraffe. In: *Mental lives,* ed., R. Campbell. Blackwell. [aGWH]
Riddoch, M. J., Humphreys, G. W., Coltheart & Funnell, E. (1988) Semantic system or systems? Neuropsychological evidence re-examined. *Cognitive Neuropsychology* 5:3–25. [aGWH, EMS]
Riddoch, M. J., Humphreys, G. W., Gannon, T., Blott, W. & Jones, V. (1999) Memories are made of this: The effects of time on stored visual knowledge in a case of visual agnosia. *Brain* 122:537–59. [aGWH]
Ritter, H. & Kohonen, T. (1989) Self-organizing semantic maps. *Biological Cybernetics* 61:241–54. [aGWH]
Rogers, T. T. & McClelland, J. L. (submitted) Semantics without categorization: A parallel distributed processing approach to the acquisition and use of natural semantic knowledge. [MALR]
Rumiati, R. I. & Humphreys G. W. (1998) Recognition by action: Dissociating visual and semantic routes to action in normal observers. *Journal of Experimental Psychology: Human Perception and Performance* 24:631–47. [EMS]
Russell, B. (1921) *Introduction to mathematical philosophy.* George Allen and Unwin. [MT]
Sacchett, C. & Humphreys, G. W. (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. [aGWH, MALR]
Saffran, E. M. (2000). A distributed model for semantic memory. *Brain and Language* 71:204–12.
Saffran, E. M., Coslett, H. B. & Keener, M. (submitted) Differences in word associations to pictorial and verbal stimuli. [EMS]
Saffran, E. M. & Schwartz, M. F. (1994) Of cabbages and things: Semantic memory from a neuropsychological point of view: A tutorial review. *Attention and Performance* 15:507–36. [aGWH, EMS]
Saffran, E. M. & Sholl, A. (1999) Clues to the functional and neural architecture of word meaning. In: *The neurocognition of language,* eds., P. Hogoort & C. Brown. Oxford University Press. [EMS]
Samson D., Pillon A. & De Wilde V. (1998) Impaired knowledge of visual and non-visual attributes in a patient with a semantic impairment for living entities: A case of a true category-specific deficit. *Neurocase* 4:273–90. [arGWH, BM]
Sartori, G. & Job, R. (1988) The oyster with four legs: A neuropsychological study on the interaction of visual and semantic information. *Cognitive Neuropsychology* 5:105–32. [arGWH]
Sartori, G., Job, R., Miozzo, M, Zago, S. & Marchiori, G. (1993a) Category-specific form-knowledge deficit in a patient with herpes simplex encephalitis. *Journal of Clinical and Experimental Neuropsychology* 15:280–99. [aGWH]
Sartori, G., Miozzo, M., & Job, R. (1993) Category-specific naming impairments? Yes. *Quarterly Journal of Experimental Psychology* 46(3):489–504.
Schiff, W., Benasich, A. & Bornstein, M. H. (1989) Infant sensitivity to audiovisually coherent events. *Psychological Research* 51:102–106. [MEA]
Schyns, P. G. (1998) Diagnostic recognition: task constraints, object information, and their interactions. In: *Object recognition in man, monkey and machine,* eds. M. J. Tarr & H. H. Bulthoff. MIT Press. [JT]
Semenza, C. & Zettin, M. (1989) Evidence from aphasia for the role of proper names as pure referring expressions. *Nature* 342:678–79. [ECL]
Shelton, J. R., Fouch, E. & Caramazza, A. (1998) The selective sparing of body part knowledge: A case study. *Neurocase* 4:339–51. [rGWH, ECL]
Sheridan, J. & Humphreys, G. W. (1993) A verbal-semantic category-specific recognition impairment. *Cognitive Neuropsychology* 10:143–84. [aGWH]
Silveri, M .C. & Gainotti, G. (1988) Interaction between vision and language in

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press

category-specific semantic impairment. *Cognitive Neuropsychology* 5:677–709. [aGWH]

Silveri, M. C., Daniele, A., Giustolisi, L. & Gainotti, G. (1991) Dissociation between knowledge of living and nonliving things in dementia of the Alzheimer type. *Neurology* 41:545–46. [aGWH]

Silveri, M. C., Gainotti, G., Perani, D., Cappelletti, J. Y., Carbone, G. & Fazio, F. (1997) Naming deficit for nonliving items: Neuropsychological and PET study. *Neuropsychologia* 35:359–67. [aGWH]

Simon, T. J. (1999) The foundations of numerical thinking in a brain without numbers. *Trends in Cognitive Sciences* 10:363–65. [MT]

Sirigu, A., Duhamel, J. R. & Poncet, M. (1991) The role of sensorimotor experience in experience in object recognition: A case of multimodal agnosia. *Brain* 114:2555–73. [aGWH]

Smith, L. B. (1989) In defense of perceptual similarity. Paper presented at the Bi-Annual meeting of the Society for Research in Child Development, Kansas City. [aGWH]

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: Human Learning and Memory* 6:174–215. [rGWH, KRL]

Snodgrass, J. G. & Yuditsky, T. (1996) Naming times for the Snodgrass and Vanderwart pictures. *Behavioural Research Methods, Instruments and Computers* 28:516–36. [arGWH]

Sober, E. (1988). *Reconstructing the past.* MIT Press. [HCB]

Solomon, G. E. A., Johnson, S. C., Zaitchik, D. & Carey, S. (1996) Like father, like son: Young children's understanding of how offspring resemble their parents. *Child Development* 67:151–71. [DZ]

Spelke, E. & Dehaene, S. (1999) Biological foundations of numerical thinking: Response to T. J. Simon [1999]. *Trends in Cognitive Sciences* 10:365–66. [MT]

Srinivas, K. (1993) Perceptual specificity in nonverbal priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 19:582–602. [MG]

Srinivas, K., Breedin, S. D., Coslett, H. B. & Saffran, E. M. (1997) Intact perceptual priming in a patient with damage to the anterior inferior temporal lobes. *Journal of Cognitive Neuroscience* 9:490–511. [aGWH, EMS]

Stewart, F., Parkin, A. J. & Hunkin, N. M. (1992) Naming impairments following recovery from herpes simplex encephalitis. *Quarterly Journal of Experimental Psychology* 44A:261–84. [aGWH]

Sutherland, N. S. (1968) Outlines of a theory of visual pattern recognition in animals and man. *Proceedings of the Royal Society of London, Series B* 171:297–317. [MG]

Takayama, Y., Sugishita, M., Akiguchi, I. & Kimura, J. (1994) Isolated acalculia due to left parietal lesion. *Archives of Neurology* 51:286–91. [MT]

Tanaka, J. W. & Taylor, M. (1991) Object categories and expertise: Is the basic level in the eye of the beholder? *Cognitive Psychology* 23:457–82. [JT]

Tarr, M. J. & Bülthoff, H. H. (1998) Image-based object recognition in man, monkey, and machine. *Cognition* 67:1–20. [MA, MG]

Tarr, M. J., Williams, P., Hayward, W. G. & Gauthier, I. (1998) Three-dimensional object recognition is viewpoint-dependent. *Nature Neuroscience* 1:275–77. [MG]

Teixeira Ferreira, C., Giusiano, B., Ceccaldi, M. & Poncet, M. (1997) Optic aphasia: Evidence of the contribution of different neural systems to object and action naming. *Cortex* 33:499–513. [EMS]

Thioux, M., Pillon, A., Samson, D., de Partz, M-P., Noël, M-P. & Seron, X. (1998) The isolation of numerals at the semantic level. *Neurocase* 4:371–89. [MT]

Thompson-Schill, S. L., Aguirre, G. K., D'Esposito, M. & Farah, M. J. (1999) A neural basis for category and modality specificity of semantic knowledge. *Neuropsychologia* 37:671–76. [arGWH, BM]

Tippett, L. J., Gloser, G. & Farah, M. J. (1996) A category-specific naming impairment after temporal lobectomy. *Neuropsychologia* 34:139–46. [aGWH]

Tranel, D., Damasio, H. & Damasio, A. (1997) A neural basis for the retrieval of conceptual knowledge. *Neuropsychologia* 35:1319–27. [aGWH]

Tranel, D., Logan, C. G., Frank, R. J. & Damasio, A.R. (1997) Explaining category-related effects in the retrieval of conceptual and lexical knowledge for concrete entities: Operationalization and analysis of factors. *Neuropsychologia* 35:1329–39. [AP]

Tversky, B. & Hemenway, K. (1984) Objects, parts, and categories. *Journal of Experimental Psychology: General* 113:169–93. [MG]

Tyler, L. & Moss, H. (1997) Functional properties of concepts: Studies of normal and brain-damaged patients. *Cognitive Neuropsychology* 14:511–45. [aGWH]

(in press) Towards a distributed account of conceptual knowledge. *Trends in Cognitive Sciences* 5:244–52. [LKT]

Tyler, L. K., Moss, H. E., Durrant-Peatfield M. & Levy, J. P. (2000) Conceptual structure and the structure of concepts. *Brain and Language* 75:195–231. [MALR, LKT]

Tyler, L. K. & Tavares, P. (in press) Is there neural specialization for pictures of living and non-living things? *Neuroimage.* [LKT]

Ullman, S. (1996) *High-level vision. Object recognition and visual cognition.* MIT Press. [MG]

Ungerleider, L. G. & Haxby, J. V. (1994) "What" and "where" in the human brain? *Current Opinion in Biology* 4:157–64. [aGWH]

Vitkovitch, M. Humphreys, G. W., & Lloyd-Jones, T. (1993) On naming a giraffe a zebra: Picture naming errors across different object categories. *Journal of Experimental Psychology: Learning, Memory and Cognition* 19:243–59. [arGWH, KRL]

Wallis, G. & Bülthoff, H. (1999). Learning to recognize objects. *Trends in Cognitive Sciences* 3:22–31.

Warrington, E. K. (1975) The selective impairment of semantic memory. *Quarterly Journal of Experimental Psychology* 27:635–57. [aGWH, MALR]

Warrington, E. K. & McCarthy, R. (1983) Category specific access dysphasia. *Brain* 106:859–78. [aGWH, BM, AP]

(1987) Categories of knowledge: further fractionations and an attempted integration. *Brain* 110:1273–96. [arGWH, ECL, BM]

(1994) Multiple meaning systems in the brain: A case for visual semantics. *Neuropsychologia* 32:1465–73. [aGWH]

Warrington, E. K. & Shallice, T. (1984) Category specific semantic impairments. *Brain* 107:829–54. [arGWH, MA, ECL, BM, MALR]

Watson, D. G. & Humphreys, G. W. (in preparation) Emergent category specificity in an interactive activation model of object recognition and naming. [aGWH]

Wierzbicka, A. (1985) *Lexicography and conceptual analysis.* Karoma. [DB]

Wu, L. L. & Barsalou, L. W. (in preparation) Grounding concepts in perceptual simulation: I. Evidence from property generation. [GSC]

Yamadori, A. & Albert, M. L. (1973) Word category aphasia. *Cortex* 9:112–25. [ECL]

Zorzi, M., Perry, C., Ziegler, J. & Coltheart, M. (1999) Category-specific deficits in a self-organizing model of the lexical-semantic system. In: *Connectionist models on cognitive neuroscience: NCPW5,* eds., D. Heinke, G. W. Humphreys & A. Olson. Springer. [aGWH]

https://doi.org/10.1017/S0140525X01294151 Published online by Cambridge University Press