The reign of typicality in semantic memory

Karalyn Patterson*

MRC Cognition and Brain Sciences Unit, Cambridge CB2 7EF, UK

This paper begins with a brief description of a theoretical framework for semantic memory, in which processing is inherently sensitive to the varying typicality of its representations. The approach is then elaborated with particular regard to evidence from semantic dementia, a disorder resulting in relatively selective deterioration of conceptual knowledge, in which cognitive performance reveals ubiquitous effects of typicality. This applies to frankly semantic tasks (like object naming), where typicality can be gauged by the extent to which an object or concept is characterized by shared features in its category. It also applies in tasks apparently requiring only access to a 'surface' representation (such as lexical decision) or translation from one surface representation to another (like reading words aloud), where typicality is defined in terms of the structure of the surface domain(s). The effects of surface-domain typicality also appear early in the time course of word and object processing by normal participants, as revealed in event-related potential studies. These results suggest that perceptual and conceptual processing form an interactive continuum rather than distinct stages, and that typicality effects reign throughout this continuum.

Keywords: perceptual versus conceptual processing; receptive versus expressive tasks; modality-specific versus modality-general representations; semantic disorders; connectionist networks; event-related potentials

Things say in one language and mean in another. No getting across that gap without the ultimate transitive, to translate.

(Powers 1992, p. 487)

1. INTRODUCTION AND THEORETICAL FRAMEWORK

Why has a paper about semantic memory come to roost in the Language section of this special issue? Although the paper could have nested in the Memory section, it also seems at home in Language, because semantic memory—although encompassing much more than language—is intimately tied to language function in humans. This link was acknowledged in Greek mythology. As one might infer from her name, Mnemosyne (daughter of Zeus and mother of the nine muses) signified memory, but she was also credited with 'giving names to everything, so that we can describe them, and converse about them without seeing them'. And Plato's *Critias* apparently urged gratitude to the memory goddess, Mnemosyne, as the source of practically all of the most important aspects of language.

Over approximately the past decade, our transatlantic research group has been developing an account of semantic memory that is framed in connectionist principles and at least parts of which are instantiated in connectionist simulations (e.g. Lambon Ralph *et al.* 2001; Plaut 2002; McClelland & Rogers 2003; Rogers &

McClelland 2004; Rogers *et al.* 2004*a*,*b*). This account was developed in part as an alternative to the proposal (e.g. by Collins & Quillian 1969) that semantic memory might be represented in terms of discrete, hierarchically organized categories of concepts and propositions about them. In this hierarchy, subordinate levels would automatically inherit the propositions applying to their superordinates, such that—having represented the fact that birds can fly—(i) one would not need to learn explicitly that owls can fly and (ii) one would need to store the fact that the ostrich, though a bird, cannot fly. As presaged by Warrington (1975) and outlined in detail by McClelland & Rogers (2003) and Rogers & McClelland (2004), the hierarchical approach encounters a variety of difficulties in accounting for the results of a body of research on conceptual knowledge, both normal and impaired. Many of these difficulties seem to be resolved in connectionist models where concepts correspond to distributed representations occupying positions in a multidimensional semantic space.

Hinton (1981) and Hinton *et al.* (1986) were the first to lay out the basic logic of capturing semantic similarity in terms of overlap among microfeatures that characterize similar concepts. This principle was then expanded by Rumelhart & Todd (1993) in their important work on the learning of semantic representations. As Rumelhart & Todd (1993) described it:

> Because learning in a distributed-representation network occurs as modification of connections among microfeatures rather than among concepts directly, generalization and transfer of learning between concepts is inescapable. When similar concepts should… be responded to similarly…the system must learn which common microfeatures similar concepts share. When important distinctions must be made between

*karalyn.patterson@mrc-cbu.cam.ac.uk

One contribution of 14 to a Discussion Meeting Issue 'Mental processes in the human brain'.

concepts that are very similar…, it must learn distinctive microfeatures that differentiate otherwise similar concepts.

(Rumelhart & Todd 1993, p. 7).

A natural consequence of this kind of system bent on discovering the important generalizations is that the most typical concepts in any category, which share the greatest number of microfeatures with other members of the category, will be robustly represented, efficiently recognized or retrieved and relatively resistant to disruption by brain disease/injury.

2. TYPICALITY FROM THE PERSPECTIVE OF NEUROPSYCHOLOGY

Semantic dementia (SD), a neurodegenerative disorder associated with atrophy of the anterior temporal lobes, is characterized by striking, and relatively selective, deterioration of semantic memory (Hodges *et al.* 1992). There are many fascinating things about the neuropsychology, neuroanatomy and neuropathology of SD, but one of the most intriguing aspects of its cognitive profile is the ubiquitous and potent impact of typicality. Two comprehensive and thoughtful early experimental studies of a few SD patients (Warrington 1975; Schwartz *et al.* 1979), conducted before the condition was given its SD label by Snowden *et al.* (1989), already provided strong hints that typicality might be an important factor. More recent research establishes just how important typicality is—not only in its impact on SD patients' degree of success/failure in virtually any cognitive domain, but also in its capacity to make sense of what can otherwise seem like inconsistent and messy data.

For an example of what typicality means here, and also an example of the way in which this variable can create order out of apparent confusion, take the case of lexical decision in SD. Lexical decision, a task created in the late 1960s/early 1970s by experimental psychologists interested in both language and memory, requires the participant to decide whether a verbal stimulus (sometimes spoken, more often written) constitutes a real word in his or her own language. CAKE is a word in English; DAKE could have been a word but is not; and the fact that people can accurately and rapidly judge the lexical status of such letter strings is considered theoretically informative. There has also been considerable controversy as to how such decisions are made. As usual in our discipline, researchers address the question of mechanism or process, at least in part, by attempting to determine which experimental manipulations affect performance. Because lexical decision is such a popular experimental task, a fair bit is now known about the variables that influence it. Some of these are straightforward and predictable: for example, people are significantly faster to make correct 'yes' judgements to high- rather than to low-frequency words, and also significantly faster to make correct judgements to words if the non-words in the experiment are not very word-like.

3. HOW DO PEOPLE, AND PATIENTS, RECOGNIZE WORDS/OBJECTS?

What happened when neuropsychologists applied the task of lexical decision to patients with a deficit of semantic memory? It appears that, with two praiseworthy exceptions (Bub *et al.* 1985; Diesfeldt 1992), these researchers tended to neglect what was already known about factors that affect normal performance and are perhaps even more likely to affect SD performance. At the time that we (Rogers *et al.* 2004*a*,*b*) were preparing to write up our first experiment on lexical decision in SD, we could find about eight publications that offered data on this topic. Now, it is true that (i) SD, like all disorders, varies in its severity (in this case, because it is a progressive disease, much of this variation is attributable to stage of decline), (ii) many of these reports were single case studies, and individual patients with the same condition—even with approximately the same degree of severity—do not always behave in the same way (though SD is a more homogeneous disorder than most), and (iii) no two of these eight studies used the same lexical decision test. Nevertheless, the outcome could be described as messy and puzzling, because this literature review told us that SD patients' success at lexical decision could vary from essentially normal, to impaired-but-above-chance, to chance level.

Our experiment (Rogers *et al.* 2004*a*,*b*), in keeping with the two notable exceptions mentioned above, demonstrated that success at lexical decision in the selfsame SD patient could vary from essentially normal to impaired-but-above-chance to chance level (or below), depending on the characteristics of the words and nonwords. We tested 22 SD patients on a two-alternative forced-choice lexical decision task, with four of these patients participating at two different stages of decline, resulting in 26 observations for each of four conditions. For a group of normal control participants, matched to the patients in age and years of education, proportions correct for the four conditions ranged from 0.99 to 0.96. By contrast—to give just two examples of individual patient data (from table 3, p. 338 of Rogers *et al.* 2004*a*,*b*)—JTh's proportions correct for the four conditions were 1.0, 1.0, 0.89 and 0.44 (the latter being chance), and the same values for patient AT were 0.94, 0.67, 0.61 and 0.22 (the latter being significantly below chance-level performance of 0.5).

What was it that varied across the conditions to yield this huge range of scores for the patients? Unsurprisingly, word frequency had a significant impact, but even more dramatically, it was the typicality of the words and non-words that determined success or failure. Typicality in this case means orthographic 'goodness' or word-likeness, as scaled by bigram and trigram frequencies. CHEESE is a nice orthographically typical word in English; SEIZE is atypical; and each of these words can be reclad in the other one's spelling pattern to yield a non-word that is either more or less orthographically typical than the real word (CHEIZE and SEESE). Asked to decide whether CHEESE or CHEIZE is the real word, almost all SD patients at any stage of severity select the correct alternative. Asked to decide between pairs like SEIZE versus SEESE, on the other hand, many patients prefer the typical non-word to the atypical real word. This preference is exacerbated when word frequency is low or stage of disease is fairly advanced or especially both,

Table 1. Eight different tasks, spanning a range of verbal versus non-verbal content and expressive versus receptive format, in which success by SD patients is significantly determined by typicality as well as frequency.

task	typicality manipulation	typical examples	atypical examples
reading aloud	spelling–sound correspondences	*sea hint*	*sew pint*
spelling to dictation	sound–spelling correspondences	*couch swerve*	*cough suave*
past-tense generation	past tense formed by +ed or change to stem?	*save (saved)* *laugh (laughed)*	*steal (stole)* *lose (lost)*
forced-choice lexical decision	real word or non-word more orthographically typical?	*cheese/cheize* *drew/driew*	*seize/seese* *view/vew*
forced-choice noun gender (Spanish)	phonological/orthographic ending of noun	*el conde* *el lacre*	*la nave* *la sierpe*
forced-choice object decision	real or non-real object more typical of object domain?	*jackal without/with a hump*	*camel with/without a hump*
delayed copy drawing	object typical of its domain?	*jackal*	*camel*
forced-choice colour decision	animal brown or not; vegetable green or not	*brown/pink bear* *green/purple asparagus*	*pink/brown flamingo* *purple/green aubergine*

such that (as indicated above for AT) performance can go significantly below chance.

This pattern, in which success is a joint function of patient severity, item familiarity and the degree to which the item in question is typical of the relevant domain, has been documented for a single group of 14 SD patients on six different tasks (Patterson *et al.* 2006). Four of the six tasks used words as stimuli: reading aloud; spelling to dictation; producing the past-tense forms of verbs from their stem (present-tense) forms; and lexical decision. The other two tasks, object decision and delayed copy drawing, used line drawings of familiar objects as stimuli. In other words: four of the six tasks were expressive or production tasks (reading aloud, spelling, past-tense generation and delayed copy), while the other two (lexical decision and object decision) were receptive. Since that study, the same pattern has been demonstrated for SD patients on two additional tasks, one non-verbal and one verbal, both receptive. The non-verbal task is selection of the correct colour for objects that have a conventionally associated colour (Adlam *et al.* 2006): in certain semantic domains, there is a typical but not universal colour (e.g. brown for animals, green for vegetables), so that one can—just as in the lexical and object decision studies described above—construct pairs in which the correctly coloured object is either more or less typical of its domain than the incorrect alternative. The verbal task is selection of the correct gender for nouns in Spanish (Sage *et al.* 2006). Some noun endings in Spanish, like -o and -a, have an almost completely predictable associated gender (masculine for -o, feminine for -a); but others are 'quasi-regular' and thus only somewhat predictable (e.g. for nouns ending in -e, 83% are masculine and 17% are feminine).

Table 1 lists these eight tasks and briefly describes the nature of the typicality manipulation in each. It should be noted that the importance of graded typicality of spelling–sound correspondences in reading aloud was first highlighted in normal performance by Glushko (1979) and in acquired surface dyslexia by Shallice *et al.* (1983). Figure 1 demonstrates the impact of three factors—patient severity, stimulus familiarity

and typicality—on SD performance in two of the language tasks, one productive (reading aloud) and the other receptive (lexical decision). Similar to the previous report described above (Rogers *et al.* 2004a,b), the SD patients' lexical decision performance from Patterson *et al.* (2006) spanned the range from nearly normal when all three of these factors were 'positive' (milder patients; higher-frequency words; words more orthographically typical than their non-word counterparts) to below chance when all three of these factors were 'negative' (more severe patients; less familiar words; words less orthographically typical than their non-word mates).

The atypical items in all of these tasks are not only vulnerable to error: as predicted, the manner in which the patients err turns the atypical stimulus into a more typical response. The two-alternative forced-choice tests (lexical decision, object decision, colour decision and Spanish noun-gender decision) cannot exactly demonstrate this point since, if the patients do make errors, they can only choose the incorrect alternatives that the test offers them; but, of course, when the incorrect alternatives offered are less typical than the correct ones, the patients make few errors. And in all of the productive tasks, the errors of commission follow the predicted pattern. Figure 2 presents the mean proportions of three types of responses, across the 14 SD patients in Patterson *et al.* (2006), in each of the four conditions of the three productive verbal tasks. Correct responses are self-explanatory. LARC errors (standing for legitimate alternative rendering of components; Patterson *et al.* 1995) are called 'legitimate' alternatives because the response would be correct if the stimulus word were typical rather than atypical. Examples of LARC errors for the three tasks are as follows: in reading aloud, *sew* → 'sue'; in spelling to dictation, 'cough' → *coff*; in past-tense generation, 'lose' → 'losed'('loozed'). 'Other errors' is a summary label for any incorrect responses that are not LARC errors. The figure demonstrates that in all tasks other errors remain fairly constant across the four conditions of the task, but as the condition shifts along from easiest (high frequency typical) to hardest (low frequency atypical), correct responses gradually decline and are replaced by LARC errors.
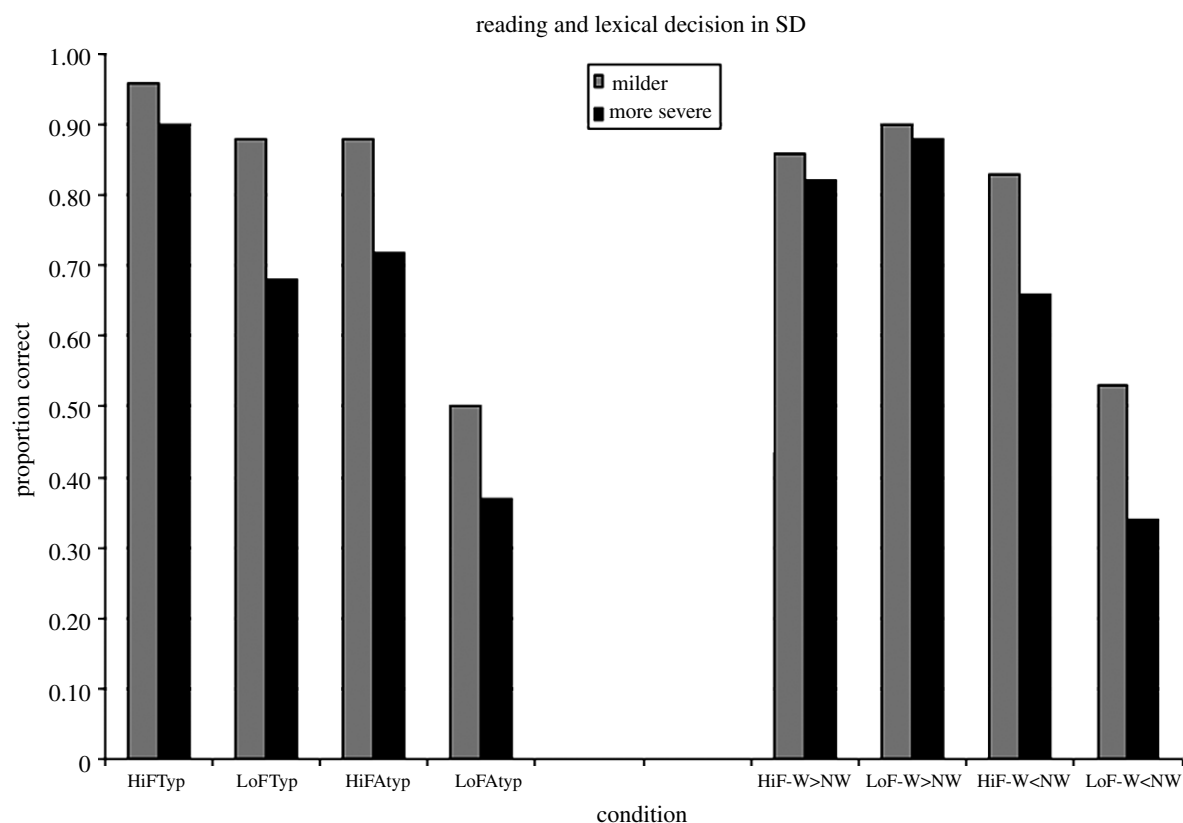
reading and lexical decision in SD

Figure 1. Results demonstrating how SD patients' success in two verbal tasks, one productive and one receptive, varies as a function of word frequency, stimulus typicality and patient severity. For both tasks, severity refers to the fact that the 14 patients have been divided into two equal subgroups on the basis of scores on a word + picture comprehension test. For the task of reading aloud (on the left), all stimuli were words, and the condition labels refer to a combination of the words' frequency (HiF, high frequency; LoF, low frequency) and their spelling–sound typicality (Typ means that the word's pronunciation is typical for its spelling and Atyp that the pronunciation is atypical for or unpredictable from the word's spelling). For the task of lexical decision (on the right), each trial consisted of a word + non-word pair. HiF and LoF in the condition labels refer to the frequency of the real word. The remaining part of each label refers to the relative orthographic typicality of the word/non-word pairs in that condition: W > NW means that the word had more typical orthographic structure than the non-word, W < NW means the reverse.

The tasks in which we specifically manipulate the typicality of stimuli and of response alternatives are of special value in revealing the impact of this variable on SD patients' success/failure rates; but, the simple tasks used commonly by neuropsychologists, such as object/picture naming or concept definitions, provide corroborating evidence. Table 2 illustrates picture naming by an SD patient in response to 25 different animals from the Snodgrass and Vanderwart picture set. Among several different points of interest that one could extract from this simple dataset are the facts: (i) that the patient's only correct responses occurred to the four most typical (also the most familiar) animals, (ii) that three of these typical animal names were assigned incorrectly to a large number of other animals—basically, the mammals—that have reasonably typical animal features, (iii) that, although all of the pictures as shown to the patient were approximately the same size, her 'choice' of a more typical incorrect label seemed somewhat appropriate to the size of the stimulus animal in real life, and, most importantly, (iv) that pictures of the atypical animals (listed at the end) were not assigned her favoured typical names. Finally, here is a single example from the concept definitions task that seems to say it all: when SD patient AM was asked 'Can you tell me what a seahorse is?', he replied 'I didn't know they had horses in the sea'.

4. THE INTERACTION BETWEEN PERCEPTUAL AND CONCEPTUAL PROCESSING

One of the messages to be derived from these findings on SD is germane to a debate about the relationship between perceptual and semantic processing. One position in this debate holds that the human brain contains structural representations or descriptions of familiar words and objects against which a stimulus can be matched to produce 'recognition'. Such structural representations are not only construed as meaning-free, but they are also thought to be a sufficient basis for recognition decisions (as in lexical or object decision) without reference to meaning (Coltheart 2004). Such a view seems to be seriously challenged by our SD data. Other findings demonstrate that SD patients succeed well at tasks requiring perceptual processing, such as matching different views of the same object (Hovius et al. 2003) or matching upper- and lower-case versions of the same letter (Cumming et al. 2006). This is as one might expect, at least on the basis that the relatively focal anterior temporal lobe atrophy in SD does not extend to the posterior regions of the brain generally thought to be crucial for perceptual processing of written words (left occipitotemporal cortex; Dehaene et al. 2002) or of seen objects (right posterior temporal cortex; Gerlach et al. 1999; Kellenbach et al. 2005). More to the point, however, our claim is that the
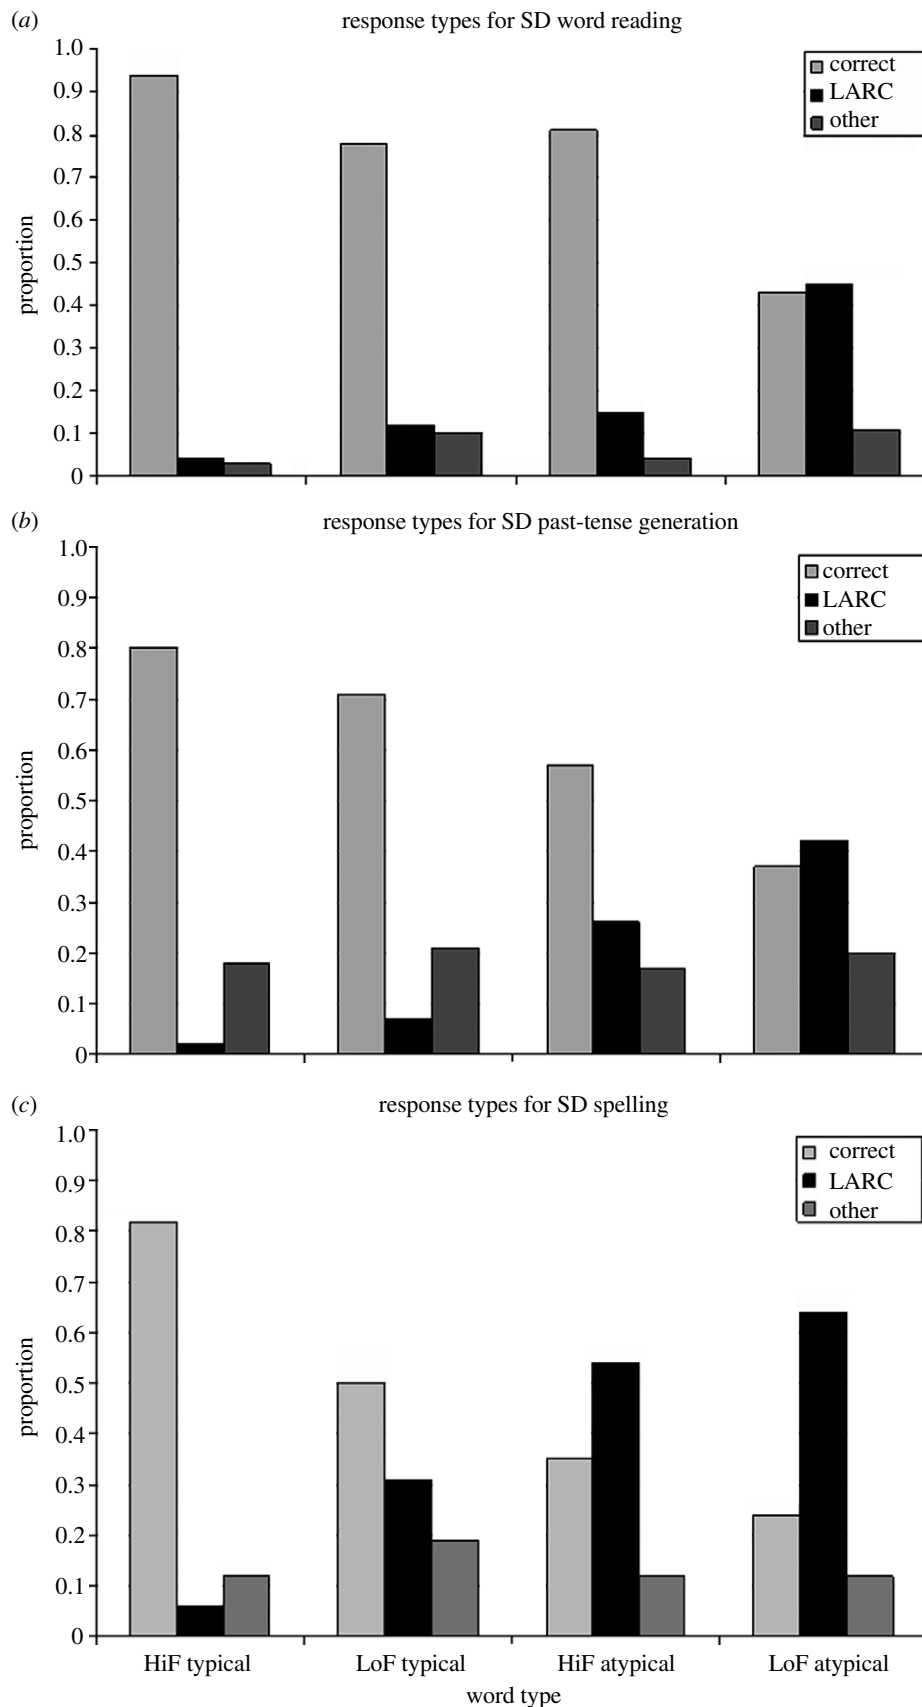
(*a*)

response types for SD word reading

(*b*)

response types for SD past-tense generation

(*c*)

response types for SD spelling

Figure 2. SD patients' performance in three productive language tasks: (*a*) reading words aloud; (*b*) generating the past tense of verbs from their stem (present-tense) forms; (*c*) spelling words to dictation. For each of the four conditions formed by crossing word frequency with domain typicality, the figure shows, across the group of 14 patients, the proportions of responses that were correct, LARC errors or some other type of error. LARC stands for legitimate alternative rendering of components, and indicates that, although the patient's response to the stimulus is incorrect, it would be correct if the stimulus were a similar one from the same domain. Such errors mainly occur to atypical items, where the response represents a more typical rendering of the stimulus (e.g. reading the word *hood* as if it rhymed with 'food'); but LARC errors can and occasionally do occur to typical items (e.g. reading the word *food* to rhyme with 'hood').

Table 2. Picture-naming responses by patient DG.

stimulus picture	naming response
dog	dog
cat	cat
horse	horse
cow	cow
pig	dog
squirrel	dog
sheep	dog
deer	dog
goat	dog
fox	dog
leopard	dog
tiger	dog
skunk	dog
racoon	dog
mouse	cat
rabbit	cat
monkey	cat
lion	horse
zebra	horse
rhino	horse
bear	horse
snake	long thing
seahorse	little thing
fish	don't know
frog	don't know

distinction between perceptual and conceptual processing is nothing like as sharp as traditional psychology textbooks might lead one to believe. As Gloor (1997, p.270) charmingly put it: '... Teuber's (1968) definition of agnosia as a normal percept stripped of its meaning, attractive as it is by its concise lapidarity, owes more to our traditional linear way of logical reasoning than to an appreciation of the underlying neurobiological mechanisms involved'.

To rephrase Gloor's (and our) position in Teuber's terms: agnosia is a percept stripped of its meaning, which is very unlikely to be a normal percept. One of our goals in studying SD has been to try to specify the nature of meaning-diminished percepts by determining what kinds of behaviours they can and cannot support. We do not claim to have a complete understanding or taxonomy with regard to this issue, but the data suggest that such meaning-diminished percepts are sufficient to cope with tasks only at the very perceptual end of this putative perceptual-to-conceptual continuum. Relevant examples here include abilities like the ones mentioned above, e.g. simultaneous matching of different views of a real object—which, after all, people can do with nonsense objects as well as real ones—or matching upper- and lower-case letters. Even these abilities turn out to be somewhat vulnerable to semantic impairment if one introduces some perceptual 'stress' or memory requirement into the tasks, such as a short filled delay between the presentations of the two different object views to be matched (Ikeda *et al.* 2006), or brief-and-masked presentation of the first of two sequentially presented different case letters to be matched (Cumming *et al.* 2006).

Perhaps the most dramatic example of how an apparently perceptual task is 'invaded' by conceptual

knowledge—meaning that performance of the task is supported by a meaningful percept in normal participants and disrupted by a meaning-diminished percept in SD patients—is delayed copy drawing. One of the abilities often taken to demonstrate relative preservation of both visuoperceptual processing and episodic memory in SD is the combination of copy and recall of the complex meaningless Rey figure. Initially, the person is asked to copy this figure with it in full view. SD patients almost invariably do this as well as age-matched controls; indeed, when we see patients in the clinic in Cambridge, impairment on copy of the Rey figure is virtually grounds on its own for rejecting a clinical diagnosis of SD. The second component of the task is that, approximately half to three-quarters of an hour later, participants are asked, without prior alerting, to reproduce the figure from memory. Normal people are not so brilliant at this kind of literal visual recall: there is a strict scoring scheme designed to capture the completeness and accuracy of the features reproduced, and control subjects' scores for delayed recall are often around 50% of the full marks achieved for direct copy. The point germane to the current discussion is that, at least in mild-to-moderate SD, not only direct copy but also delayed recall of the Rey complex figure tends to be normal: although the patients' scores are sometimes at the lower end of the normal range, they are certainly not several standard deviations below it as is consistently observed even in very early Alzheimer's disease. If SD patients achieve reasonably normal success in both copying and recalling meaningless figures, one might predict that they would do the same with line drawings of real, familiar objects. This prediction is, however, wrong or at least half wrong: the patients' ability to copy is fine but their recall is not.

We do this experimental task in a somewhat different manner to the procedure for administering the Rey figure. In one test session, we ask the patient to copy line drawings of objects, such as a rhinoceros or a watering can, with the stimulus picture present. On a different day, we show them each picture and let them study it for perhaps 5 s so that they have a good idea of what it looks like; they are not asked to name it, nor does the experimenter name it for them. The drawing is then removed from view and the patient counts aloud from 1 to 15, which takes patients (and controls) only approximately 10 s. The patient is then asked to draw what he or she was looking at before counting (again, no name mentioned). Figure 3, an example from Bozeat *et al.* (2003), illustrates both an SD patient's copy of the rhinoceros and his recall of it a mere 10 s later. The direct reproduction would perhaps not win any awards, but it qualifies as a thoroughly recognizable facsimile. The delayed copy, on the other hand, has lost virtually all of its unique features and looks like a much more generic animal shape, perhaps a pig or a long-snouted dog.

How can one explain the discrepancy, both between SD patients' direct and delayed copy of meaningful objects and between their delayed recall of meaningless and meaningful figures? Our proposal is that the two tasks that they perform with essentially normal accuracy are at the very perceptual end of the perceptual↔conceptual continuum. Direct copy of a
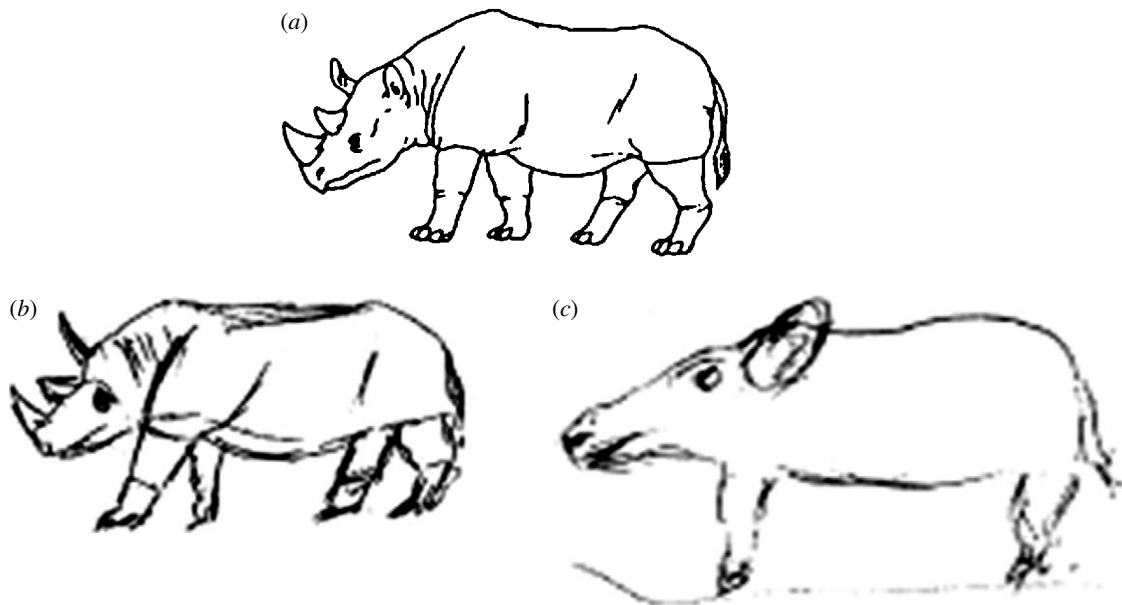
Figure 3. Drawings by an SD patient (DS) of (*a*) the stimulus picture of a rhinoceros: (*b*) the drawing is DS's copy of the picture with it present and (*c*) the drawing is his copy of the picture approximately 10 s after it had been removed from view.

meaningful figure, like simultaneous matching of two different views of an object, can be done without reference to conceptual knowledge, as witnessed by the fact that both normal people and patients can do these tasks on nonsense figures. Delayed recall of meaningless figures can be done (though not all that skilfully by either patients or controls) without reference to conceptual knowledge because it has to be: there are no obvious concepts to refer to in performing this task. But delayed recall of a meaningful drawing—which would presumably be as difficult as recall of a meaningless figure if it were being performed solely on the basis of literal stimulus memory—can invoke, indeed probably cannot avoid invoking, conceptual knowledge. Normal participants, in delayed copy drawing of a rhinoceros, do not need to rely on literal visual memory in order to draw its horns: once they have recognized the stimulus picture as a rhino, semantic memory insists that it must have horns. Our hypothesis is that exactly the same reliance on semantic knowledge occurs, cannot fail to occur, in SD patients; but, because their conceptual knowledge is significantly degraded, they have only recognized the rhino as some kind of generic animal, and that is what their delayed reproduction represents. This is agnosia as redefined by Gloor (1997) and by us: a percept that is stripped of its meaning and is accordingly abnormal.

The hornless, armourless rhinoceros in figure 3 may seem like a particularly striking example, and, of course, not all of the SD patients' delayed copy drawings are quite so unfaithful to their stimuli; however, there are many other examples like the rhino, including humpless camels and four-legged ducks (Lambon Ralph & Howard 2000; Bozeat *et al.* 2003). What, apart from patient severity, modulates success when SD patients do this task? Our old friend typicality. Rogers *et al.* (2004*a*,*b*) established that the patients are likely both to omit distinctive properties of objects that should have such unusual features, and to intrude common properties when they attempt to recall objects that happen not to possess such shared features.

Rhinoceros horns and camel humps are, of course, atypical, distinctive features, and while no birds have four legs, the most typical animals (of the cat–dog–horse ilk) all do. If SD patients' conceptual recognition of a duck amounts basically to 'animal', then it is not all that surprising that, once the duck has disappeared, they assign it four legs. Such errors in the patients' drawings could be considered a non-verbal parallel of the LARC errors described above by which the patients render words more typical in reading aloud, spelling or past-tense inflection.

5. WHERE DOES TYPICALITY HAVE ITS IMPACT? EVIDENCE FROM BRAIN IMAGING

Patients with SD have a relatively selective disorder of conceptual knowledge, and we have in the past argued that—even in tasks seeming to require minimal access to semantic information, such as reading single words aloud or delayed copy drawing—all deficits observed in SD stem from the central semantic degradation (Rogers *et al.* 2004*a*,*b*; Patterson *et al.* 2006; Woollams *et al.* in press). The sections above were designed to illustrate the pervasive impact of typicality on performance in SD, in both more and less obviously semantic tasks (e.g. object naming and reading aloud, respectively). If we attribute all of the patients' deficits to the central semantic disorder, then it would seem logical to assume that the typicality effect derives from the structure of semantic memory, as indeed we have in the past also argued (McClelland *et al.* 1995; Rogers *et al.* 2004*a*,*b*). On the other hand, if there is no sharp distinction between perceptual and conceptual processing of the meaningful things (objects and words) that people encounter every minute of their waking lives, is it correct to attribute the typicality effect entirely to the conceptual end of processing?

Functional imaging in normal individuals might be helpful in answering this question, since it would allow the researcher to ask where/when typicality manipulations affect patterns of brain activity. There is as yet no

evidence from a technique with high spatial resolution, like functional magnetic resonance imaging (fMRI), on where in the brain typicality effects might be observed. There is, however, some evidence from two event-related potential (ERP) studies (Hauk *et al.* 2006, in press); these were motivated by our lexical decision and object decision experiments on typicality effects in SD patients and, like them, independently manipulated the typicality and the lexicality or authenticity of the stimuli. Apart from the neurological normality of the participants and the recording of ERPs, these studies differed from the SD experiments primarily in using a yes/no rather than a forced-choice procedure. Presentation of each word and non-word (or real and non-real object) as a separate event rather than in pairs was necessary to measure the time course of ERP signals, and estimate their brain sources, separately for the different experimental conditions.

In the lexical decision study (Hauk *et al.* 2006), there was a very early significant effect of typicality, at approximately 100 ms post-stimulus onset, where both words and non-words with atypical orthographic patterns (e.g. *yacht*, *cacht*) elicited stronger activation than words and non-words with typical spelling patterns (e.g. *yart*, *cart*). Acknowledging that there are always limits on the precision with which one can determine the sources of ERP effects, our procedure for source estimation (the 'minimum norm solution'; see Hauk 2004 for details) highlighted relatively focal activation in the left mid-posterior temporal lobe in association with the early typicality effect. In other words, although *yacht* is a familiar word and *yart* is not, at this stage of the word recognition process and in this region of the brain, *yart* appears to be processed more efficiently. This left mid-posterior temporal area may correspond to the region that, in fMRI experiments, responds significantly more to the written forms of both words and pseudowords than to letter strings (like a series of consonants) that are not possible words (Dehaene *et al.* 2002; McCandliss *et al.* 2003). Of major interest here are the facts that this ERP typicality effect: (i) occurred so soon following stimulus presentation and (ii) was 'blind' to lexicality. A significant main effect of lexicality, in the form of a stronger ERP response to pseudowords than real words, was observed later at 200 + ms. These facts suggest that substantial sensitivity to stimulus typicality is characteristic of early perceptual processing as well as later semantic processing. The one remaining ERP finding of particular salience in Hauk *et al.* (2006) is that, about midway between the typicality effect at approximately 100 ms and the lexicality effect at approximately 200 ms post-stimulus, there was a significant interaction of these two variables, with stronger activation for atypical than typical words, and the source estimation for this effect was in the left anterior temporal cortex, in or near the location of major atrophy in SD. This finding is compatible with, although it does not in any sense confirm, the hypothesis that the convergence centre of the semantic system works harder on atypical than typical familiar things before deciding that they are real/meaningful.

Our ERP study on object decision (Hauk *et al.* in press) yielded a strikingly similar time course of ERP

effects, although the difference between real and non-real objects (the equivalent of the lexicality effect in lexical decision, which we have termed an authenticity effect for object decision) occurred much later, at approximately 480 ms. This difference can probably be explained in terms of the far greater visual complexity of line drawings of objects relative to printed words. Similar to the lexical decision study, we obtained an early effect, at 116 ms post-stimulus, of typicality: atypical line drawings, whether authentic or not, produced greater activation than typical drawings, and source estimation localized this effect to bilateral occipitotemporal cortex. Once again, in other words, the evidence suggests that typicality has a pervasive influence that is early and perceptual as well as later and semantic.

6. CONCLUDING COMMENTS

Some other things, that there was insufficient time/ space to discuss here, are already known about the impact of typicality on both normal and abnormal cognition. But §5 of this little treatise, concerning the impact of typicality on more perceptual versus more conceptual processing, is a good example of how much remains to be understood about the nature of typicality and the manner in which it affects representations of different kinds of information, at different locations in the brain. What is gradually becoming clear, however, is: (i) how pervasive the influence of typicality seems to be and (ii) how well-suited connectionist models—like Rogers *et al.* (2004*a,b*)—are to capturing and explaining such effects.

REFERENCES

Adlam, A.-L. R., Patterson, K., Rogers, T. T., Nestor, P., Salmond, C. H., Acosta-Cabronero, J. & Hodges, J. R. 2006 Semantic dementia and fluent primary progressive aphasia: two sides of the same coin? *Brain* **129**, 3066–3080. (doi:10.1093/brain/awl285)

Bozeat, S., Lambon Ralph, M. A., Graham, K. S., Patterson, K., Wilkin, H., Rowland, J., Rogers, T. T. & Hodges, J. R. 2003 A duck with four legs: Investigating the structure of conceptual knowledge using picture drawing in semantic dementia. *Cogn. Neuropsychol.* **20**, 27–47. (doi:10.1080/02643290244000176)

Bub, D., Cancelliere, A. & Kertesz, A. 1985 Whole-word and analytic translation of spelling to sound in a nonsemantic reader. In *Surface dyslexia* (eds K. Patterson, J. C. Marshall & M. Coltheart), pp. 15–34. Hillsdale, NJ: Erlbaum.

Collins, A. M. & Quillian, M. R. 1969 Retrieval time from semantic memory. *J. Verb. Learn. Verb. Behav.* **8**, 240–247. (doi:10.1016/S0022-5371(69)80069-1)

Coltheart, M. 2004 Are there lexicons? *Q. J. Exp. Psychol. A: Hum. Exp. Psychol.* **57**, 1153–1171.

Cumming, T., Patterson, K., Verfaellie, M. & Graham, K. S. 2006 One bird with two stones: letter-by-letter reading in pure alexia and semantic dementia. *Cogn. Neuropsychol.* **23**, 1130–1161. (doi:10.1080/02643290600674143)

Dehaene, S., Le Clec'H, G., Poline, J. B., Le Bihan, D. & Cohen, L. 2002 The visual word form area: a prelexical representation of visual words in the fusiform gyrus. *Neuroreport* **13**, 321–325. (doi:10.1097/00001756-200203040-00015)

Diesfeldt, H. F. A. 1992 Impaired and preserved semantic memory functioning in dementia. In *Memory functioning in dementia* (ed. L. Backman). Amsterdam, The Netherlands: Elsevier.

Gerlach, C., Law, I., Gade, A. & Paulson, O. B. 1999 Perceptual differentiation and category effects in normal object recognition: a PET study. *Brain* **122**, 2159–2170. (doi:10.1093/brain/122.11.2159)

Gloor, P. 1997 *The temporal lobe and limbic system*. New York, NY; Oxford, UK: OUP.

Glushko, R. J. 1979 The organization and activation of lexical knowledge in reading aloud. *J. Exp. Psychol. Hum. Percept. Perform.* **5**, 674–691. (doi:10.1037/0096-1523.5.4.674)

Hauk, O. 2004 Keep it simple: a case for using classical minimum norm estimation in the analysis of EEG and MEG data. *Neuroimage* **21**, 1612–1621. (doi:10.1016/j.neuroimage.2003.12.018)

Hauk, O., Patterson, K., Woollams, A., Pulvermuller, F., Watling, L. & Rogers, T. T. 2006 [Q]: When would you prefer a SOSSAGE to a SAUSAGE? [A]: At about 100 ms. ERP correlates of orthographic typicality and lexicality in written word recognition. *J. Cogn. Neurosci.* **18**, 818–832. (doi:10.1162/jocn.2006.18.5.818)

Hauk, O., Patterson, K., Woollams, A., Pye, E., Pulvermuller, F. & Rogers, T. T. In press. How the camel lost its hump: the impact of typicality on ERP signals in object decision. *J. Cogn. Neurosci.*

Hinton, G. E. 1981 Implementing semantic networks in parallel hardware. In *Parallel models of associative memory* (eds G. E. Hinton & J. A. Anderson), pp. 161–188. Hillsdale, NJ: Erlbaum.

Hinton, G. E, McClelland, J. L. & Rumelhart, D. E. 1986 Distributed representations. In *Parallel models of associative memory*, vol. 1 (eds D. E. Rumelhart, J. L. McClelland & the PDP Research Group). Cambridge MA: MIT Press.

Hodges, J. R., Patterson, K., Oxbury, S. & Funnell, E. 1992 Semantic dementia: progressive fluent aphasia with temporal lobe atrophy. *Brain* **115**, 1783–1806. (doi:10.1093/brain/115.6.1783)

Hovius, M., Kellenbach, M. L., Graham, K. S., Hodges, J. R. & Patterson, K. 2003 What does the object decision task measure? Reflections on the basis of evidence from semantic dementia. *Neuropsychology* **17**, 100–107. (doi:10.1037/0894-4105.17.1.100)

Ikeda, M., Patterson, K., Graham, K. S., Lambon Ralph, M. A. & Hodges, J. R. 2006 "A horse of a different colour": do patients with semantic dementia recognise different versions of the same object as the same? *Neuropsychologia* **44**, 566–575. (doi:10.1016/j.neuropsychologia.2005.07.006)

Kellenbach, M. L., Hovius, M. & Patterson, K. 2005 A PET study of visual and semantic knowledge about objects. *Cortex* **41**, 107–118.

Lambon Ralph, M. A. & Howard, D. 2000 Gogi aphasia or semantic dementia? Simulating and assessing poor verbal comprehension in a case of progressive fluent aphasia. *Cogn. Neuropsychol.* **17**, 437–465. (doi:10.1080/02643290410784)

Lambon Ralph, M. A., McClelland, J. L., Patterson, K., Galton, C. J. & Hodges, J. R. 2001 No right to speak? The relationship between object naming and semantic

impairment: neuropsychological evidence and a computational model. *J. Cogn. Neurosci.* **13**, 341–356. (doi:10.1162/08989290151137395)

McCandliss, B. D., Cohen, L. & Dehaene, S. 2003 The visual word form area: expertise for reading in the fusiform gyrus. *Trends Cogn. Sci.* **7**, 293–299. (doi:10.1016/S1364-6613(03)00134-7)

McClelland, J. L. & Rogers, T. T. 2003 The parallel distributed processing approach to semantic cognition. *Nat. Rev. Neurosci.* **4**, 310–322. (doi:10.1038/nrn1076)

McClelland, J. L., McNaughton, B. L. & O'Reilly, R. C. 1995 Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychol. Rev.* **102**, 419–457. (doi:10.1037/0033-295X.102.3.419)

Patterson, K., Suzuki, T., Wydell, T. & Sasanuma, S. 1995 Progressive aphasia and surface alexia in Japanese. *Neurocase* **1**, 155–165. (doi:10.1093/neucas/1.2.155)

Patterson, K., Lambon Ralph, M. A., Jefferies, E., Woollams, A., Jones, R., Hodges, J. R. & Rogers, T. T. 2006 'Pre-semantic' cognition in semantic dementia: six deficits in search of an explanation. *J. Cogn. Neurosci.* **18**, 169–183. (doi:10.1162/jocn.2006.18.2.169)

Plaut, D. C. 2002 Graded modality-specific specialization in semantics: a computational account of optic aphasia. *Cog. Neuropsychol.* **19**, 603–639. (doi:10.1080/02643290244000112)

Powers, R. 1992 *The Gold Bug Variations*. London, UK: Harper Perennial.

Rogers, T. T. & McClelland, J. L. 2004 *Semantic cognition: a parallel distributed processing approach*. Cambridge, MA: MIT Press.

Rogers, T. T., Lambon Ralph, M. A., Garrard, P., Bozeat, S., McClelland, J. L., Hodges, J. R. & Patterson, K. 2004a The structure and deterioration of semantic memory: a neuropsychological and computational investigation. *Psychol. Rev.* **111**, 205–235. (doi:10.1037/0033-295X.111.1.205)

Rogers, T. T., Lambon Ralph, M. A., Hodges, J. R. & Patterson, K. 2004b Natural selection: the impact of semantic impairment on lexical and object decision. *Cogn. Neuropsychol.* **21**, 331–352. (doi:10.1080/02643290342000366)

Rumelhart, D. E. & Todd, P. M. 1993 Learning and connectionist representations. In *Attention and performance XIV: synergies in experimental psychology, artificial intelligence, and cognitive neuroscience* (eds D. E. Meyer & S. Kornblum), pp. 3–30. Cambridge, MA: MIT Press.

Sage, K., Heredia, C. G., Berthier, M. & Lambon Ralph, M. A. 2006 The impact of semantic impairment on lexical gender in Spanish. Paper presented to the British Neuropsychological Society, October.

Schwartz, M. F., Marin, O. S. & Saffran, E. M. 1979 Dissociations of language function in dementia: a case study. *Brain Lang.* **7**, 277–306. (doi:10.1016/0093-934X(79)90024-5)

Shallice, T., Warrington, E. K. & McCarthy, R. 1983 Reading without semantics. *Q. J. Exp. Psychol.* **35A**, 111–138.

Snowden, J. S., Goulding, P. J. & Neary, D. 1989 Semantic dementia: a form of circumscribed temporal atrophy. *Behav. Neurol.* **2**, 167–182.

Warrington, E. K. 1975 Selective impairment of semantic memory. *Q. J. Exp. Psychol.* **27**, 635–657.

Woollams, A. M., Lambon Ralph, M. A., Plaut, D. C. & Patterson, K. In press. SD-Squared: on the association between semantic dementia and surface dyslexia. *Psychol. Rev.*