Exemplar Models and Category-Specific Deficits

Koen Lamberts

University of Warwick

Laura Shapiro

University of Birmingham

Mailing address:

Koen Lamberts

Department of Psychology

University of Warwick

Coventry CV4 7AL

United Kingdom

e-mail: K.Lamberts@warwick.ac.uk

Exemplar Models and Category-Specific Deficits

In recent years, there have been numerous reports of patients with brain damage who show selective identification or recognition deficits for objects from specific categories (see Forde, in press; Humphreys & Forde, 2000, for reviews). The most common deficit appears to be a selective impairment in the identification of living things, accompanied by relatively unimpaired recognition or identification of artificial or non-living objects. However, despite the large number of reported cases with category-specific processing deficits, there is still no agreement on the mechanisms that produce these deficits. It is not even clear whether all such cases can be understood in terms of a single process or mechanism, or whether category-specific deficits can be caused by a variety of different factors. In this chapter, we explore category-specific deficits from a theoretical viewpoint that evolved from recent research on perceptual categorization and identification. Although some efforts have been made to model category-specific deficits with connectionist models (e.g., Farah & McClelland, 1991; Humphreys, Lamote, & Lloyd-Jones, 1995), we are not aware of any attempts to apply classical models of categorization and identification[1] to the neuropsychological data on category-specificity (with the exception of a study by Dixon, Bub, & Arguin, 1997, which will be discussed in detail later).

Current theories of categorization (and identification, which is a special case in which each object forms its own category) can be divided into five groups. The first group is that of exemplar models, which assume that categorization of an object depends on the similarity of that object to instances in memory (e.g., Estes, 1994; Kruschke, 1992; Medin & Schaffer, 1978; Nosofsky, 1986). Second, there are decision-bound models, which are based on the multidimensional generalisation of classical signal detection theory (Ashby & Lee, 1991; Ashby & Maddox, 1993). According to these models, stimuli correspond to points in a multidimensional space. The perceptual representations of stimuli are assumed to be variable from trial to trial, due to intrinsic noise in the perceptual system. For categorization, people are assumed to establish linear or non-linear category decision bounds in the multidimensional stimulus space. Categorization depends on the position of the stimulus representation on a given trial relative to the decision bounds. Third are the models which explain category decisions based on the application of formal rules (e.g., Nosofsky, Palmeri, & McKinley, 1994). The fourth group contains various connectionist models of categorization (e.g., Gluck & Bower, 1988), which attempt to explain categorization in terms of associative links between input information and response alternatives. Finally, there have been recent proposals for combined models of categorization, which integrate elements from two or more of the theories listed above (e.g., Ashby, Alfonso-Reese, Turken, & Waldron, 1998; Erickson & Kruschke, 1998).

In this chapter, we focus exclusively on exemplar models of categorization and identification. There are several reasons for this choice. Exemplar models have an impressive empirical track record. They can explain categorization and identification of a wide range of different stimuli in a wide range of situations (e.g., Lamberts, 1994, 1995, 1998, 2000; Lamberts & Freeman, 1999; Medin & Schaffer, 1978; Nosofsky, 1984, 1986, 1987, 1991-a, 1992; Nosofsky & Palmeri, 1997). Conceptually, exemplar models are well developed and understood, and their relations with other classes of models have been explored in great detail (e.g., Alfonso-Reese & Ashby, 1995; Ashby & Maddox, 1993; Nosofsky, 1991-b). Finally, exemplar models provide a unifying framework for a broad range of seemingly disparate cognitive tasks (e.g., Brockdorff & Lamberts, 2000; Estes, 1994; Lamberts, in press; Nosofsky, 1991-a; Nosofsky & Zaki, 1998; Palmeri, 1997).

Exemplar models of categorization and identification

According to exemplar models, learning involves the storage of instances in memory. Exemplar models do not assume that learning involves the computation of summary representations for categories or other groups of stimuli (as presumed in prototype models or rule-based models). Instead, it is assumed that each encounter with a stimulus leaves a separate trace in memory, and that subsequent categorization, identification or recognition depends on the retrieval of these specific memory traces. There are usually no constraints on the kind of information that can be contained in a memory trace. Exemplar information can be perceptual (referring to structural or surface properties of the object; Humphreys, Riddoch, & Quinlan, 1988) or semantic (referring to aspects of its meaning).

Probably the most successful exemplar model to date is Nosofsky's (1986) Generalized Context Model (GCM). The GCM assumes that stimuli can be defined as points in a multidimensional psychological space. Each dimension of the space corresponds to a particular aspect of the stimulus (such as colour, size, etc.). Although the GCM is intended primarily as a model of perceptual categorization, dimensions can also refer to abstract or semantic stimulus attributes. Similarity between stimuli is defined as a decreasing function of the distance between the stimuli in the psychological space. We will use the following definition of similarity (which is a special case of the similarity notion of the GCM, see Nosofsky, 1986):

$$\eta_{ij} = \exp[-c(\sum_{p=1}^{P} w_p | x_{ip} - x_{jp} |)] . \qquad (1)$$

In this equation, the similarity between the representations of two stimuli (i and j) is a decreasing exponential function of the weighted sum of differences between the stimuli along the stimulus dimensions. w_p is the weight of dimension p, and x_{ip} and x_{jp} are the values of stimulus i and stimulus j on dimension p. If a dimension is more heavily weighted, a difference along that dimension will affect the similarity value more than a difference along a dimension with less weight. The parameter c is an index of discriminability. This index determines how quickly similarity decreases as a function of the distance between the stimulus representations (see Lamberts, 1994, for an extensive discussion of the role of this parameter). If c is high, stimuli are highly discriminable, meaning that even a small difference between them will result in a relatively low similarity value. Unless we explicitly note otherwise, we will simply omit dimension weights from this equation in our applications of the GCM (thus assuming that all dimensions have the same weight).

If there are two alternative categories, the GCM assumes that the probability that a stimulus is classified in a given category depends on the summed similarity of that stimulus to the exemplars of the category on the one hand, and the total similarity of the stimulus to all exemplars in both categories on the other hand. The version of the GCM that we will use states that the probability that a stimulus i is classified into category C is given by

$$P(R_C \mid S_i) = \frac{\sum_{j \in C} \eta_{ij}}{\sum_k \eta_{ik}} . \qquad (2)$$

The GCM also applies to identification tasks, in which there are as many response alternatives as stimuli (in other words, each stimulus requires a unique response). The only difference with the categorization model is in the choice rule, which becomes

$$P(R_j \mid S_i) = \frac{\eta_{ij}}{\sum_k \eta_{ik}}. \qquad (3)$$

This rule states that the probability of a response j to stimulus i is a function of the similarity between stimulus i and exemplar j (which has associated response j) on the one hand, and the total similarity of stimulus i to all exemplars in memory on the other hand. Because self-similarity is 1 in the GCM, the probability of correct identification thus becomes

$$P(Correct \mid S_i) = \frac{1}{\sum_k \eta_{ik}}.$$

Although exemplar models have been used primarily to account for categorization, identification and recognition in normal individuals, there have been a few attempts to apply exemplar models to neuropsychological data. An important application by Dixon, Bub and Arguin (1997) will be discussed later in this chapter. Nosofsky and Zaki (1998) have recently shown that an exemplar model can provide insight into complex patterns of performance in patients with amnesia. The starting point of their work was a series of experiments by Knowlton and Squire (1993), in which groups of normal and amnesic patients categorised or made old/new judgements for sets of visual patterns. The results showed that the normal controls performed much better than the patients on old/new recognition, whereas both groups performed at a similar level in categorization. Knowlton and Squire (1993) interpreted this result in terms of multiple memory systems, with an implicit system responsible for the acquisition of categorical knowledge, and a declarative system responsible for old/new recognition. In amnesics, the declarative system was supposed to be damaged (causing poor recognition performance), but an intact implicit system would still allow normal categorization performance. Nosofsky and Zaki (1998) showed that it was not necessary to assume that the dissociation in task performance reflected an underlying dissociation in processing systems. Specifically, Nosofsky and Zaki (1998) demonstrated that a single exemplar model explained Knowlton and Squire's (1993) data, if it was assumed that brain damage led to a parameter change. A model in which the discriminability parameter $\underline{c}$ (see Equation 1) had a smaller value for amnesic patients than for normal controls produced the dissociation between categorization and recognition observed by Knowlton and Squire (1993), without having to assume separate subsystems for these tasks. Moreover, the model also explained the results from two other studies (Knowlton, Mangels, & Squire, 1996; Knowlton, Squire, & Gluck, 1994), on exactly the same assumptions. The discriminability parameter ($\underline{c}$ in Equation 1) determines how steeply similarity decreases with increasing distance between stimulus representations. If $\underline{c}$ is high, processing is very selective, and perfect matches between representations are weighted far more in decision making than imperfect matches. If $\underline{c}$ is low, even poor matches produce relatively high similarity values. The value of $\underline{c}$ can have a great impact on the behaviour of exemplar models. For instance, Lamberts (1994) showed that changes in discriminability can produce model behaviour that ranges from a nearest-neighbour model (in which only the most similar exemplar determines categorization) to a nearly-linear prototype model (in which similarity to the "average" or prototype of the category determines choice). Nosofsky and Zaki's (1998) results show the potential of exemplar models to provide a single-systems explanation for dissociations that seem to invite a multiple-systems interpretation. The importance of this work for the interpretation of category-specific deficits is obvious. In the past, these deficits too have been explained in terms of multiple semantic or memory systems (e.g., Caramazza & Shelton, 1998; Sartori & Job, 1988; Silveri & Gainotti, 1988; Silveri, Daniele, Giustolisi, & Gainotti, 1991), and we will attempt to show that an exemplar account produces category specificity, without having to assume multiple storage or retrieval systems.

Exemplar storage and category-specific deficits

Now that we have defined the principles of exemplar models, we can explore the implications of such models for understanding the category-specific deficits reported in the literature. Although most deficit studies have used identification as the main task, we will discuss both identification and categorization.

An important aspect of any model of neuropsychological deficits is the implementation of brain damage in terms of the model's components and processes. Because we do not know the physiological mechanisms that might support processes such as those defined in exemplar models, we can only postulate plausible ways in which neurological damage could alter the characteristics of the psychological model. In this chapter, we will investigate two possible effects of neurological damage. Following Nosofsky and Zaki (1998), we will explore the effects of decreased stimulus discriminability. We will also investigate the nonselective loss of features of stored exemplars. Not only are these two plausible consequences of brain damage, their effects are also quite similar to those resulting from other possible damage processes (such as loss of the ability to process particular dimensions). Of course, it is possible that neurological damage has other, unanticipated effects. It is conceivable, for instance, that a loss of exemplar information would be accompanied by noise in decision making. For the purpose of clarifying and exploring the predictions of exemplar theories for category-specific deficits, assuming that brain damage results in random loss of features or decreased discriminability is quite sufficient.

Before we could start the modelling work with the GCM, we had to explore the structure of the objects within the categories that we were going to study. The stimulus dimensions that may underlie the representations of living and non-living objects are unknown, so we had to make apriori assumptions about the structure of the living and non-living categories. These assumptions are crucial for the modelling work. Even without formal demonstration, it is obvious that category-specific deficits are unlikely to emerge through feature loss (or any other mechanism) if the damage is non-selective and the categories have the same underlying structure. On these conditions, category-specific deficits would only occur if the damage somehow affected exemplars from one category much more than exemplars from the other category. If damage is non-selective (as we will assume in all the modelling), one would expect both categories to suffer to the same extent. Category-specific deficits could only emerge exceptionally, as a result of random variation in the damage effects. However, for categories that are fairly large and that contain objects with a large number of features (such as the living and non-living categories), category-specific deficits would be extremely rare. Moreover, both categories would have the same likelihood of being selectively affected, and this is contradicted by the far higher incidence of category-specific deficits for living things in patients with brain damage. From these considerations, it is quite clear that exemplar models will only predict systematic category-specific deficits as a consequence of non-selective damage if the categories involved are somehow different from each other. In the following section, we explore the differences between living and non-living objects that may be relevant for understanding category-specific deficits.

Differences between living and non-living object categories

In many studies in which category-specific deficits were reported, the stimuli were taken from Snodgrass and Vanderwart's (1980) collection of object drawings. Snodgrass and Vanderwart (1980) have supplied norms for name agreement, image agreement, familiarity and visual complexity for their picture set, so it is relatively straightforward to determine whether the living and non-living objects used in studies that report category-specific deficits differed on these variables, or to design studies in which these variables are controlled. Interestingly, when familiarity, word frequency and name agreement were matched for living and non-living stimuli, Funnell and Sheridan (1992) found that a disproportionate impairment for living things disappeared in one patient. Gaffan and Heywood (1993) and Stewart, Parkin and Hunkin (1992) also found that the poorer performance for living things compared to non-living things observed in their patient disappeared once word frequency, familiarity and visual complexity were matched for the two categories. However, Farah, Meyer and McMullen (1996) found that when two of their patients were tested on exactly the same set of pictures but with further replications, their selective deficits for living things remained. Gainotti and Silveri (1996) and Kurbat (1997) also found that category specific effects occurred in their patients when the normed variables were controlled. Together, these results indicate that category-specific effects are not purely due to any differences in familiarity, word frequency and name agreement that might exist between the categories of living and non-living things.

A potentially far more relevant difference between living and non-living object categories concerns the similarity relations that exist between the category members. Various studies have suggested that the similarity structures of the living and non-living categories in the Snodgrass and Vanderwart (1980) set are not equivalent (e.g. Humphreys, Riddoch, & Quinlan, 1988; Gaffan and Heywood, 1993; Humphreys, Lamote and Lloyd-Jones, 1995). By similarity, we mean similarity in the purely perceptual sense; for instance, pictures of a horse and a dog are perceptually similar because they both contain the same components (such as head, neck, body and legs). When Humphreys et al. (1988) asked normal participants to list the parts of living and non-living things, living things showed up as having more shared parts than non-living things. The authors also compared the outline contours of standardised drawings from different categories by normalising all the Snodgrass and Vanderwart pictures for size and orientation, and then overlaying each picture with every other picture from the same category on a grid and calculating the overlap. The living things tended to have higher degrees of contour overlap than the non-living things.

Further evidence that perceptual similarity is higher within living categories than within non-living categories was provided by Gaffan and Heywood (1993), using the Snodgrass and Vanderwart (1980) picture set. Normal subjects made more errors naming living than non-living things when stimulus quality was degraded, indicating that living things are less visually discriminable (i.e., more perceptually similar) than non-living things. The authors also trained monkeys to make discriminative responses to pictures of living and non-living things and found that the monkeys took longer to learn the living responses than non-living responses. Specifically, their difficulty in distinguishing among living things increased steeply as the number of living things to be discriminated increased. The authors concluded that the high levels of perceptual overlap in living categories caused the difficulties in discrimination between these items.

We have replicated these findings ourselves. Ten subjects (undergraduate students) gave pairwise ratings of perceptual similarity for a randomly chosen set of 15 living and 15 non-living pictures from the Snodgrass and Vanderwart set. The subjects were asked to ignore what the pictures actually represented and to concentrate entirely on the perceptual characteristics of the drawings on the screen. They were asked to give each pair a rating between 0 and 9, where 0 was for a pair that look nothing like each other and 9 was for a pair that was almost identical. The living pairs were rated as being more similar to each other than the non-living pairs; the mean rating for living pairs was 3.69, compared to 2.49 for the non-living pairs. This effect was significant, $t(9) = 10.39$, $p < .001$.

We have also achieved the same result by using the reaction time to decide that two pictures were different as a measure of similarity. It was assumed that for two very dissimilar pictures, participants would be able to decide very quickly that they were different. However, for two very similar pictures, more features would have to be processed before the differences became apparent and reaction times would be much slower (see Lamberts & Brockdorff, 2000). In this experiment we used 48 of the pictures selected by Funnell and Sheridan (1992), where norm values for word frequency, familiarity and visual complexity were equivalent for the living and non-living pictures. This particular set is useful to judge whether the living and non-living categories differ in their similarity profile, even when these other factors are controlled. The 24 living pictures and 24 non-living pictures were grouped into pairs of identical and different pictures within the living and non-living categories. The mean time taken to correctly decide that 2 pictures were different was significantly longer for the living pairs (531 ms) than for the non-living pairs (518 ms), $F(1, 19) = 21.79$, $p < .001$, indicating that the living pictures are more similar to each other than the non-living pictures.

We have outlined good evidence that perceptual similarity is not equivalent within living and non-living categories. All the above findings indicate that for the Snodgrass and Vanderwart (1980) set, the living pictures are perceptually more similar to each other than the non-living pictures. The Snodgrass and Vanderwart set is used to test most patients with category-specific deficits. Therefore, we decided to design the categories used in the simulation work with the GCM in such a way that one category contained elements that were more similar to each other than the other category.

Simulating category-specific deficits

In all the modelling work that we report in this chapter, two categories of simulated objects were used. Each category contained 20 exemplars, and each exemplar consisted of 15 continuous dimensions. The categories were constructed in the following way. First, we defined a prototype for each category. The prototype of the first category had a value of 0 on all 15 dimensions, whereas the prototype of the second category had a value of 1 on all dimensions. Next, the prototypes were used to generate the exemplars within each category (the prototypes themselves were not part of the categories). Exemplars were generated by random distortion of the category prototypes. Each exemplar from the first category was obtained by adding a random number to the prototype value for each dimension. The random numbers for this category were drawn from a rectangular distribution with a mean of 0 and range from –0.3 to +0.3. The exemplars from the second category were generated in a similar fashion, except that the range of the random numbers was from –1 to +1. As a result, there was more variability in the second category, and the exemplars from this category were less similar to each other than the exemplars from the other category. We will call the first category the “homogeneous” category, and the second the “heterogeneous” category. Within each category, a total of 20 exemplars were generated.

The effects of brain damage were first simulated by randomly removing features from the 40 exemplars that made up the memory set. The expected proportion of deleted features varied between 0 (intact memory) and 0.9 (severe loss of feature information). At each level of damage, we simulated 1000 cases. In the simulation, the dimension(s) that corresponded to a missing exemplar feature were simply omitted from the similarity calculations. For each case, the GCM was applied to generate a predicted proportion of correct identification responses across all the exemplars within each category. The model was also used to predict proportions of correct categorization responses across all exemplars within each category. The only model parameter that needed clamping was $\underline{c}$, the discriminability index (see Equation 1). To obtain a better overview of the model’s range of predictions, we repeated the entire simulation experiment with different values of $\underline{c}$.

Figure 1 shows the simulation results, separately for identification and categorization within the two categories, and for two different values of c (3.0 and 5.0, respectively).

Insert Figure 1 about here

The simulated results for the identification task showed a strong contrast between the two categories. Loss of features had a stronger detrimental effect on the identification of the objects in the homogeneous category than on identification of the objects in the heterogeneous category. For the categorization task, however, the opposite pattern occurred. Categorization of the heterogeneous objects declined more rapidly than categorization of the homogeneous objects. (Note that the absolute difference in performance between the identification and categorization tasks is partly due to the different levels of expected chance performance in the two tasks). This pattern occurred for both values of c in this simulation, and further modelling work showed that it occurred across a wide range of category structures and parameter settings. Whenever one category contained exemplars that were more similar to each other than the exemplars in the other category, feature loss affected identification performance more in the category with similar exemplars, while it affected categorization more in the category with relatively dissimilar exemplars. Intuitively, it is easy to understand why this pattern emerges. Identification, in which each stimulus requires a unique response, is more difficult if a stimulus has close neighbours. If features are lost, a stimulus may become less distinguishable from one or more other stimuli, and performance will drop. For categorization, however, close neighbours help, because the probability of a correct response depends on the total similarity of a stimulus to all exemplars within a category. If a stimulus has many close neighbours, even a large proportion of features can be lost before performance drops significantly.

Figure 2 summarises the model's predictions for identification and categorization of homogeneous and heterogeneous objects, on the assumption that feature memory is intact, but for different levels of stimulus-exemplar discriminability (c). It is immediately clear that the effects of variation in c are almost identical to the effects of random feature loss. This is not surprising, given that random feature loss will reduce the average distance between stimuli, just as lower discriminability does.

Insert Figure 2 about here

The simulation results show that a difference in internal similarity structure is sufficient to explain category-specific identification deficits as a consequence of non-selective damage. At the same time, the model also shows that an identification deficit for a homogeneous category should be accompanied by a categorization deficit for a heterogeneous category, if the similarity structure of these two categories is the main factor responsible for selective deficits. In the following sections, we explore whether reported patient data are compatible with these predictions.

Identification of living and non-living objects

Case studies of category-specific deficits have concentrated mainly on identification performance. In almost all reported cases, identification is disproportionately poor for stimuli from living categories (e.g., Basso, Capitani, &

Laiacona, 1988; Caramazza & Shelton,1998; De Renzi & Lucchelli, 1994; Farah & Wallace, 1992; Forde, Francis, Riddoch, Rumiati, & Humphreys, 1997; Sartori & Job, 1984; Sheridan & Humphreys, 1993; Silveri & Gainotti, 1988; Warrington & Shallice, 1984). For example, Warrington and Shallice (1984) reported 4 patients with difficulties in visual identification. The comprehension capacities of 2 patients (J.B.R and S.B.Y.) were examined in detail and a dissociation between living and non-living things was observed. J.B.R. correctly identified 6% of living pictures compared to 90% of non-living pictures. A similar pattern of performance was observed for S.B.Y., who correctly identified none of the living pictures and 75% of the non-living pictures. Similarly, Farah and Wallace (1992) reported a patient, T.U., whose naming was disproportionately poor for fruits and vegetables even when familiarity and name frequency were taken into account. In naming the Snodgrass and Vanderwart line drawings, he correctly named 54% of fruits and vegetables compared with 87% of other categories. His naming latencies were also much slower for fruits and vegetables than for other categories. Caramazza and Shelton (1998) also reported a patient, E.W., with a disproportionate impairment in naming living things. For a subset of the Snodgrass and Vanderwart pictures, matched for familiarity and name frequency, E.W. correctly named 55% of animals and 82% of non-animals.

In addition to recording absolute identification performance for different categories of objects, it is also informative to look at the different types of errors that are reported for living and non-living things. Arguin, Bub, and Dudek (1996; see also Dixon, Bub, & Arguin, 1997) reported the case of a patient, E.L.M., who showed a selective impairment for naming living objects (39% correct responses to 66 pictures of animals, birds, insects, fruits, and vegetables), with relatively intact naming of non-living objects (88% correct responses to 79 pictures of tools, clothing, instruments, etc.). Interestingly, Arguin et al. reported a confusion matrix from a word-picture matching task, which showed that ELM tended to confuse the identities of living objects with similar shapes (e.g., banana, carrot and cucumber were often confused with each other, and also apple, onion, orange and tomato). This response pattern is entirely in agreement with the GCM’s predictions. According to the model, the probability that two individual items are confused depends directly on their similarity, and the model therefore predicts clusters of confusion between similar items.

Several other reports have shown that patients with category-specific deficits for living things are more likely to confuse living things with other living things than non-living things with other non-living things (e.g., Moss, Tyler, Durrant-Peatfield, & Bunn, 1998; Stewart, Parkin, & Hunkin, 1992; Warrington & Shallice, 1984). For instance, Stewart, Parkin and Hunkin’s (1992) patient, H.O., gave the name of another object in the same category for 31.6% of the living things that were shown to him, whereas only 9.1% of his errors for non-living things were in the same category. Moss et al. (1998) tested their patient RC in a word-picture matching task, in which a spoken word was presented and the patient had to select the corresponding picture from an array of four. In addition to the target, there was always one distracter from the same category, and two distracters from other categories. R.C. was significantly more accurate in identifying non-living targets than living targets. On living targets, the vast majority of his errors (86%) were within-category errors, in which he chose an alternative from the same category as the target.

According to the GCM, the probability that an incorrect response will be given that corresponds to a stimulus from the same category as the target (which we will call a within-category identification error) is given by

$$P(Error\ within\ category \mid S_i) = \frac{(\sum_{j \in C} \eta_{ij}) - 1}{\sum_k \eta_{ik}}$$

if the target stimulus belongs to category C and all symbols have the same meaning as before. This probability is exactly the same as the probability of categorization of the stimulus into the correct category, minus the probability of identifying the stimulus correctly. The probability of an error outside category for a stimulus i from category 1 is

$$P(Error\ outside\ category \mid S_i) = \frac{\sum_{j \in C_{alt}} \eta_{ij}}{\sum_k \eta_{ik}}$$

which is identical to the probability that the stimulus would be categorised into the wrong category. Figure 3 shows the probabilities of identification errors within and outside category across the homogeneous and heterogeneous stimuli used in the simulation, for different proportions of feature loss.

Insert Figure 3 about here

Within-category identification errors are much more likely than outside-category errors for the homogeneous stimuli. For the heterogeneous stimuli, within-category identification errors are still more likely than outside-category errors, but the difference is much smaller than for the homogeneous stimuli. These predictions are in agreement with the confusion data that have been discussed.

Arguin, Bub, and Dudek (1996) and Dixon, Bub and Arguin (1997) have reported another series of identification experiments with their patient E.L.M., which are very relevant to our account of category-specific deficits. As we noted before, an important problem for any similarity-based account of category-specific deficits is that the dimensions which are used to encode real objects are unknown. In mainstream categorization research, this problem is overcome by using artificial stimuli that vary along well-defined dimensions (e.g., Lamberts, 1995, 1998; Lamberts & Freeman, 1999; Medin & Schaffer, 1978), or by using techniques such as multi-dimensional scaling to infer the locations of stimuli with continuous dimensions in psychological space (e.g., Nosofsky, 1992; Nosofsky & Palmeri, 1997). Arguin et al. (1996) chose the first solution, and overcame the problem of unknown stimulus dimensions by using a set of artificial stimuli that varied on three dimensions. The stimuli were computer-generated blobs that varied in bending, elongation and tapering. Four different blobs were simultaneously presented in the four corners of the screen, for a short time. One of these blobs was then centrally presented, and E.L.M. was asked to point to that blob's former location. The most important manipulation in this experiment was the structure of the set of four blobs presented in a single trial. In single-dimension sets, the four blobs differed on one dimension only, and all had the same values on the other two dimensions. In conjunction sets, the blobs varied on two dimensions, while the third dimension was held constant. The conjunction sets were designed such that both variable dimensions needed to be processed in order to identify the stimulus. E.L.M. performed consistently better on the single-dimension sets (29% errors) than on the conjunction sets (57% errors). Arguin et al. further showed that E.L.M. did not have a perceptual problem with processing multiple dimensions, from which they inferred that his problems with conjunctive sets were based in memory.

Dixon et al. (1997) argued that E.L.M.'s performance in the Arguin et al. (1996) study was incompatible with any of the existing neuropsychological theories of category-specific deficits. However, they did show that an exemplar model of identification (Kruschke's, 1992, ALCOVE) could explain the results. ALCOVE is based directly on the GCM. It is presented as a connectionist model, in which backpropagation is used to learn dimension weights and exemplar strengths. Because Dixon et al. (1997) did not attempt to fit learning curves, there is no real need to apply ALCOVE to their data, and we will therefore present their arguments in terms of the GCM. Dixon et al. (1997) point out that optimal performance in the single-dimension task would be achieved by selectively weighting the relevant dimension more than the two irrelevant dimensions. This is predicted by the attention-optimisation hypothesis (see Lamberts, 1999; Nosofsky, 1986), which states that subjects will tend to use dimension weights that maximise performance in a given task. ALCOVE is designed explicitly to implement this selective weighing process. In the conjunctive conditions, the same weighting mechanism should emphasise the two relevant dimensions. Dixon et al. (1997) show that the exemplar model predicts E.L.M.'s performance, if it is assumed that dimension weights are close to optimal and stimulus discriminability is low (which is the same assumption as that of Nosofsky & Zaki, 1998). Arguin et al.'s (1996) results are thus entirely compatible with our exemplar-based account. The exemplar model also explains other, potentially more puzzling aspects of E.L.M.'s performance. In another series of experiments, Dixon et al. (1997) showed that the dimensionality effect of the earlier study was modulated by semantics. When the same shapes were paired with semantically close or disparate sounds or labels, E.L.M.'s error rates in the conjunctive task showed a strong correlation with the semantic proximity of these sounds or labels, whereas there was no such relation in the single-dimension task. An exemplar model predicts these effects, if it is assumed that discriminability is low, that optimal weighting occurs, and that the stimulus labels form an intrinsic part of the stimulus representation, such that similarity depends both on visual and semantic features (Dixon et al., 1997).

Categorization of living and non-living objects

Although most studies of category-specific deficits have focused on naming or identification, it is important for an evaluation of our account to contrast identification performance with categorization. Objects can be categorised at many different levels (see Murphy & Lassaline, 1997), and we are not aware of many systematic comparisons between categorization performance at different levels in patients with category-specific identification deficits. However, there have been several studies that show preserved categorization abilities in categories for which naming deficits occurred, as predicted by the GCM.

Forde, Francis, Riddoch, Rumiati, & Humphreys (1997) carried out a number of experiments with their patient S.R.B., and found that his naming ability (tested with the Snodgrass and Vanderwart pictures, photographs and real objects) was impaired more for living objects than for non-living objects. Reaction times were slower and more errors were made for items from living categories and this was not confounded by name-frequency, familiarity or visual complexity. Forde et al. (1997, Experiment 19) also examined S.R.B.'s ability to categorise living and non-living 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 only errors were classifying a watermelon as a vegetable and an artichoke as a fruit.

Caramazza and Shelton (1998) observed that their patient, E.W., made many more errors when naming living pictures than non-living pictures. The authors commented that the nature of the errors E.W. made for living things were quite different to the kinds of errors made for non-living things. In particular, for 34/47 living pictures, E.W. either said “I have no idea what it is” or produced a semantically related response. For example, when shown a picture of a Zebra, E.W. responded, “Gorilla, I think but I’m not sure”. By contrast, she only produced 5 of 137 semantic or “don’t know” responses to items in other semantic categories. Caramazza and Shelton (1998) also observed that E.W. could distinguish animals from artefacts so she had no selective impairment in categorising animals. E.W. was also shown to have no difficulty in answering questions concerning attributes shared by all members of a category. This indicates, again, that her problem lay in distinguishing amongst highly similar exemplars whereas she was unimpaired for tasks that require grouping.

Moss, Tyler, Durrant-Peatfield and Bunn (1998) have looked explicitly at categorization versus identification performance for their patient, R.C. Tested with the Snodgrass and Vanderwart picture set, R.C. was able to name 50% of the pictures of artefacts, compared to only 9% of pictures of living things. Similar results were obtained in a naming task with a different set of stimuli (photographs matched for familiarity). 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 non-living things. It is interesting to note that Stewart, Parkin and Hunkin's (1992) patient H.O. produced a similar pattern of errors in a naming task. For 10.5% of the errors H.O. made to living pictures, he gave the superordinate name but did this for only 4.5% of his errors for non-living pictures. This indicates that, for the living things, he was sometimes aware of the category the object was from even when he could not identify the object. Moss et al. (1998) also carried out a direct test of R.C.'s ability to categorise colour photographs of living and non-living 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 non-living objects (83% correct) was below the range for controls.

To summarise, we have shown that for patients, identification is most often worse for living things than for non-living things, whereas classification performance shows the opposite trend. Patients can often categorise living things even when the individual name is not known. They are often able to identify the superordinate for living things and often confuse the target with a member of the same category. This is less often the case for non-living things. In the cases where grouping of living and non-living things is compared, patients' selective deficit for living things disappears and in some cases, they perform better for living than non-living things for these tasks. The GCM predicts all these differences.

Is a similarity-based account sufficient to explain category-specific deficits?

Thus far, we have demonstrated that a classical exemplar model of categorization and identification predicts significant aspects of category-specific deficits in patients with visual agnosia, if it is assumed that the categories of living and non-living objects have a different internal similarity structure. The question remains whether such a simple account is sufficient to explain all aspects of performance in patients with category-specific deficits.

The model that we proposed certainly fails to explain why some patients show identification deficits for non-living objects. Indeed, although the vast majority of studies have reported identification deficits for living things, there have been a few reports of patients with impaired identification of non-living objects (e.g., Hillis & Caramazza, 1991; Sacchett & Humphreys, 1992; Warrington & McCarthy, 1983, 1987, 1994). Without attempting to dismiss these findings, it is worthwhile to explore their implications for our account. In the three case studies by Warrington and McCarthy (1983, 1987, 1994), the stimuli were not matched for a number of potential confounds, including frequency, visual complexity and familiarity (see Funnell & Sheridan, 1992). It is possible that one or more of these variables contributed to the unusual outcome of the case studies. However, the same argument does not apply to the results reported by Hillis and Caramazza (1991) and Sacchett and Humphreys (1992). Hillis and Caramazza (1991) used the same stimuli and test procedures with two patients, one of whom showed a selective deficit for living things, whereas the other was impaired on non-living things. Sacchett and Humphreys (1992) controlled for a number of confounding variables, and still observed a selective naming deficit for non-living things.

It is clear that the exemplar model that we used cannot explain these last two sets of data, without making additional assumptions about the structure of categories of living and nonliving objects or about the effects of brain damage. In fact, the model can readily produce a double dissociation like the one reported by Hillis and

Caramazza (1991) (or any other dissociation, for that matter) if it is assumed that damage is selective and somehow affects exemplars from one category more than exemplars from another category. Alternatively, one could assume that particular stimulus dimensions would have different weights in the identification of living and nonliving objects. A commonly cited distinction is that between functional and perceptual features, where the former are assumed to be more important for the identification of nonliving objects, and the latter have more weight in the identification of living things (see Farah & McClelland, 1991; Sacchett & Humphreys, 1992; Warrington & Shallice, 1984). On these assumptions, the model can trivially produce a double dissociation between identification of living and non-living objects by selective damage to the representation or the processing of perceptual or functional stimulus dimensions. However, neither of these accounts are very satisfactory. Apart from being largely ad hoc, they fail to explain why selective deficits for non-living objects are so rare. Perhaps it is safest to reserve judgement about the importance of deficits for non-living objects, until more cases have been documented and the crucial variables that underlie these deficits are better understood.

Other data that are potentially challenging for the model are those obtained in conditions where similarity within categories has been controlled or measured, and in which category-specific identification deficits appear unrelated to similarity differences between categories. For instance, Sartori, Miozzo and Job (1993) claim that higher perceptual similarity between living things is not the cause of their patient’s impairment. Sartori et al. (1993) tested their patient, Michelangelo, using drawings of animals and artefacts taken from the Snodgrass and Vanderwart set plus line drawings in a similar style. 7 subsets of animals and 6 sub-sets of artefacts were chosen with high within-set visual and semantic similarity. Undergraduates rated the global similarity of items in each subset and the ratings for animal and artefact subsets did not differ significantly. When Michelangelo was asked to name these pictures, his selective deficit for living things remained.

These results were not confirmed in a number of other studies, in which similarity-related effects have been observed directly. Livingstone (1988) studied a patient's ability to point to a named picture either amongst visually similar or visually dissimilar distracters, and found that the patient performed much better when in the visually dissimilar condition, even within living categories for which he was impaired.

Forde, Francis, Riddoch, Rumiati and Humphreys (1997) provided direct evidence that similarity rather than the living/non-living distinction was the crucial variable in S.R.B.’s performance. The authors have taken into account the structural similarity of the pictures when testing S.R.B.’s naming ability. They used the 76 Snodgrass and Vanderwart pictures from Humphreys et al. (1988), who grouped these pictures into structurally similar categories (animals, fruit, vegetables) and structurally different categories (clothing, tools, furniture). Structural similarity was determined by the number of rated common parts per category and the average percentage of contour overlap relative to other objects from the same category. Forde et al. found that S.R.B. was significantly more impaired at naming items from structurally similar categories (71% correct) compared to structurally dissimilar categories (95% correct). In fact, when a regression analysis was carried out on S.R.B.’s reaction time to name 59 Snodgrass and Vanderwart pictures, the authors found that the living /non-living distinction was not a significant predictor of performance when measures of structural similarity were taken into account. Instead, degree of contour overlap with other category members was the only significant predictor. Further evidence that structural similarity is behind S.R.B.'s deficit is that his subordinate naming of items from 2 categories of particular interest to him, dogs and cars, was very poor. He was worse for dogs (17%) than cars (57%) but this general problem at retrieving subordinate names is consistent with the hypothesis that high levels of structural similarity between category members make identification difficult.

Together with the results that we have reviewed in the previous sections, such findings suggest that similarity is often at the root of patients' problems in identifying living things. Studies in which a selective deficit remains after similarity is controlled run counter to the general trend. In the case of Sartori et al.'s (1993) study, it is doubtful as to whether their measure of similarity was adequate. They found that when ratings of overall similarity were obtained for sub-sets of pictures, no differences between living and non-living sets were found. This certainly does not stand up to findings from our own experiments, in which more rigorous measures of similarity were made for pictures from the Snodgrass and Vanderwart set. Humphreys et al.'s (1998) analyses of shared parts and contour overlap, Gaffan and Heywood's (1993) discriminability analyses, and our own pairwise ratings and reaction time experiments all showed the living pictures to be reliably more similar to each other than the non-living pictures. The stimuli used by Sartori et al. (1993) mainly included Snodgrass and Vanderwart pictures, so it is likely that the average perceptual similarity still differed between categories. Sartori et al.'s (1993) assessment of similarity was perhaps not sensitive enough to show relevant differences between the categories.

Conclusions

We have demonstrated in this chapter that a classical exemplar model of categorization and identification explains many aspects of category-specific deficits in patients with brain damage, on the assumption that the categories involved have a different similarity structure. The most important prediction from the model concerns the apparent dissociation between identification and categorization. The neuropsychological data that allow a comparison between these two tasks generally support the 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.

Another important topic for further work would be to explore the relation between our proposal and other single-systems accounts of category-specific deficits. For instance, there is a complex relation between various connectionist models of category-specific deficits (e.g., Devlin et al., 1998; McRae et al., 1997; Rogers & Plaut, this volume) and our exemplar account. Several connectionist accounts rely heavily on the notion that patterns of correlations between features are important for understanding category-specific deficits, and there is independent empirical evidence to confirm the importance of these correlation patterns (e.g., McRae, this volume). Exemplar models preserve complete information about feature correlations, and they seem therefore excellently suited to explain the role of correlations in category-specific deficits.

References

Alfonso-Reese, L. A., & Ashby, F. G. (1995). Categorization as probability-density estimation. Journal of Mathematical Psychology, 39, 216-233.

Arguin, M., Bub, D. & Dudek, G. (1996). Shape integration for visual object recognition and its implication in category specific visual agnosia. Visual Cognition, 3(3), 221-275.

Ashby, F. G., Alfonso-Reese, L. A., Turken, A.U., & Waldron, E. M. (1998). A neuropsychological theory of multiple systems in category learning. Psychological Review, 105, 442-481.

Ashby, F. G., & Lee, W. W. (1991). Predicting similarity and categorization from identification. Journal of Experimental Psychology: General, 120, 150-172.

Ashby, F. G., & Maddox, T. W. (1993). Relations between prototype, exemplar and decision bound models of categorization. Journal of Mathematical Psychology, 37, 372-400.

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.

Brockdorff, N., & Lamberts, K. (2000). A feature-sampling account of the time course of old-new recognition judgments. Journal of Experimental Psychology: Learning, Memory and Cognition, 26, 77-102.

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

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

Devlin, J., Gonnerman, L., Anderson, E. & Seidenberg, M. (1998). Category-specific deficits in focal and widespread damage: A computational account. Journal of cognitive Neuroscience, 10:1, 77-94.

Dixon, M.J., 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, 1085-1130.

Erickson, M. A., & Kruschke, J. K. (1998). Rules and exemplars in category learning. Journal of Experimental Psychology: General, 127, 107-140.

Estes, W. K. (1994). Classification and cognition. New York: Oxford University Press.

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

Farah, M. J., 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, 137-154.

Farah, M.J. & Wallace, M.A. (1992). Semantically bounded anomia: Implications for the neural implementation of naming. Neuropsychologia, 30(7), 609-621.

Forde, E.M.E. (in press). Category specific recognition impairments. In G.W. Humphreys (Ed.), Case Studies in the Neuropsychology of Vision. Psychology Press.

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), 403-458.

Funnell, E. & Sheridan, J. (1992). Categories of knowledge? Unfamiliar aspects of living and nonliving things. Cognitive Neuropsychology, 9, 135-153.

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, 118-128.

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, 357-389.

Gluck, M. A. & Bower, G. H. (1988). From conditioning to category learning: An adaptive network model. Journal of Experimental Psychology: General, 117, 225-244.

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

Humphreys, G. W., & Forde, E. M. E. (2000). Category-specific deficits: A review and presentation of the Hierarchical Interactive Theory (HIT). Manuscript submitted for publication.

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, 535-586.

Humphreys, G.W., Riddoch, M.J. & Quinlan, P.T. (1988). Cascade processes in picture identification. Cognitive Neuropsychology, 5, 67-103.

Knowlton, B. J., Mangels, J. A., & Squire, L. R. (1996). A neostriatal habit learning system in humans. Science, 273, 1399-1402.

Knowlton, B. J., & Squire, L. R. (1993). The learning of categories: Parallel brain systems for item memory and category knowledge. Science, 262, 1747-1749.

Knowlton, B. J., Squire, L. R., & Gluck, M. A. (1994). Probabilistic classification learning in amnesia. Learning and Memory, 1, 106-120.

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

Kurbat, M.A. (1997). Can the recognition of living things really be selectively impaired? Neuropsychologia, 35(6), 813-827.

Lamberts, K. (1994). Flexible tuning of similarity in exemplar-based categorization. Journal of Experimental Psychology: Learning, Memory and Cognition, 20, 1003-1021.

Lamberts, K. (1995). Categorization under time pressure. Journal of Experimental Psychology: General, 124, 161-180.

Lamberts, K. (1998). The time course of categorization. Journal of Experimental Psychology: Learning, Memory, and Cognition, 24, 695-711.

Lamberts, K. (1999). Attention supports perceptual categorization. Visual Cognition, 6, 93-99.

Lamberts, K. (2000). Information-accumulation theory of speeded categorization. Psychological Review, 107.

Lamberts, K. (in press). Feature-sampling models of categorization and recognition. Quarterly Journal of Experimental Psychology.

Lamberts, K., & Brockdorff, N. (2000). A unified model of matching and recognition. Manuscript submitted for publication.

Lamberts, K., & Freeman, R. P. J. (1999). Building object representations from parts: Tests of a stochastic sampling model. Journal of Experimental Psychology: Human Perception and Performance, 25, 904-926.

Livingstone, M.S. (1988). Art, illusion and the visual system. Scientific American, 258, 78-86.

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

Medin, D. L. and Schaffer, M. M. (1978). Context theory of classification learning. Psychological Review, 85, 207-238.

Moss, H. E., Tyler, L.K., Durrant-Peatfield, M.R. & Bunn, E. M. (1998). 'Two eyes of a see-through': Impaired and intact semantic knowledge in a case of a selective deficit for living things. Neurocase, 4, 291-310.

Murphy, G. L., & Lassaline, M. E. (1997). Hierarchical structure in concepts and the basic level of categorization. In K. Lamberts & D. Shanks (Eds.), Knowledge, concepts and categories (pp. 93-131). Hove, UK: Psychology Press.

Nosofsky, R. M. (1984). Choice, similarity, and the context theory of classification. Journal of Experimental Psychology: Learning, Memory and Cognition, 10, 104-114.

Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. Journal of Experimental Psychology: General, 115, 39-57.

Nosofsky, R. M. (1987). Attention and learning processes in the identification and categorization of integral stimuli. Journal of Experimental Psychology: Learning, Memory, and Cognition, 13, 87-109.

Nosofsky, R. M. (1991-a). Tests of an exemplar model for relating perceptual classification and recognition memory. Journal of Experimental Psychology: Human Perception and Performance, 17, 3-27.

Nosofsky, R. M. (1991-b). Relation between the rational model and the context model of categorization. Psychological Science, 2, 416-421.

Nosofsky, R. M. (1992). Similarity scaling and cognitive process models. Annual Review of Psychology, 43, 25-53.

Nosofsky, R. M., & Palmeri, T. J. (1997). An exemplar-based random walk model of speeded classification. Psychological Review, 104, 266-300.

Nosofsky, R. M., Palmeri, T. J., & McKinley, S. C. (1994). Rule-plus-exception model of classification learning. Psychological Review, 101, 53-79.

Nosofsky, R. M., & Zaki, S. R. (1998). Dissociations between categorization and recognition in amnesic and normal individuals: An exemplar-based interpretation. Psychological Science, 9, 247-255.

Palmeri, T. J. (1997). Exemplar similarity and the development of automaticity. Journal of Experimental Psychology: Learning, Memory, and Cognition, 23, 324-354.

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.

Sartori, G. & Job, R. (1988). The oyster with four legs: A neuropsychological study on the interaction of visual and semantic information. Cognitive Neuropsychology, 5(1), 105-132.

Sartori, G., Miozzo, M. & Job, R. (1993). Category-specific naming impairments? Yes. Quarterly Journal of Experimental Psychology, 46(3), 489-504.

Sheridan, J. & Humphreys, G.W. (1993). A verbal-semantic category-specific recognition impairment. Cognitive Neuropsychology, 10 (2), 143-184.

Silveri, M.C. & Gainotti, G. (1988). Interaction between vision and language in category-specific semantic impairment. Cognitive Neuropsychology, 5(6), 677-709.

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, 545-546.

Snodgrass, J. G., & Vanderwart, M. (1980). A standardised 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.

Stewart, F., Parkin, A. J. & Hunkin, N. M. (1992). Naming impairments following recovery from herpes simplex encephalitis. Quarterly Journal of Experimental Psychology, 44, 261-284.

Warrington, E. K. & McCarthy, R. (1983). Category-specific access dysphasia. Brain, 106, 859-878.

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

Warrington, E. K. & McCarthy, R. (1994). Multiple meaning systems in the brain: A case for visual semantics. Neuropsychologia, 32, 1465-1473.

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

Footnotes

[1] In this chapter, we use the terms categorization, identification and recognition in the following way. Categorization refers to a decision situation, in which objects have to be assigned to categories. The number of categories (i.e., the number of response alternatives) is smaller than the number of different objects that can occur, which implies that several objects require the same response. Identification is a special case of categorization, in which each object forms its own category. In an identification task, the number of response alternatives is the same as the number of possible objects. Finally, recognition refers specifically to old-new recognition, in which a decision is made as to whether an object has been encountered before, regardless of its category membership or identity.

Author notes

The work in this chapter was supported by a BBSRC grant to K. L., and a BBSRC Research Studentship to L. S. We thank Andrew Olson, Evan Heit, Nick Chater and Noellie Brockdorff for their helpful comments.

Figure 1

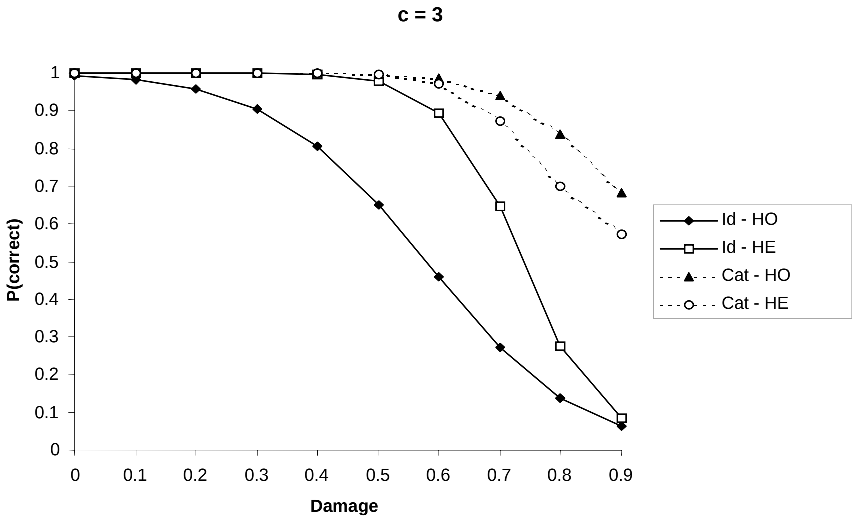
c = 3
P(correct)
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
Damage
Id - HO
Id - HE
Cat - HO
Cat - HE

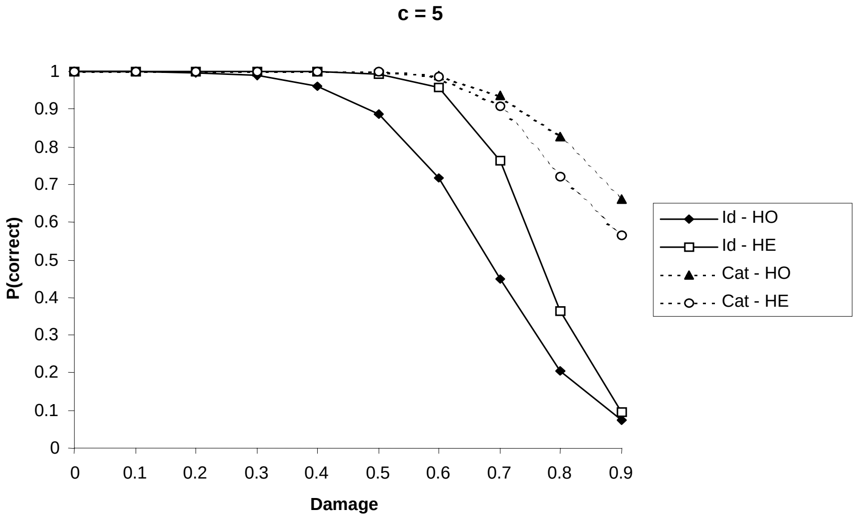
c = 5
P(correct)
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
Damage
Id - HO
Id - HE
Cat - HO
Cat - HE

Figure 2

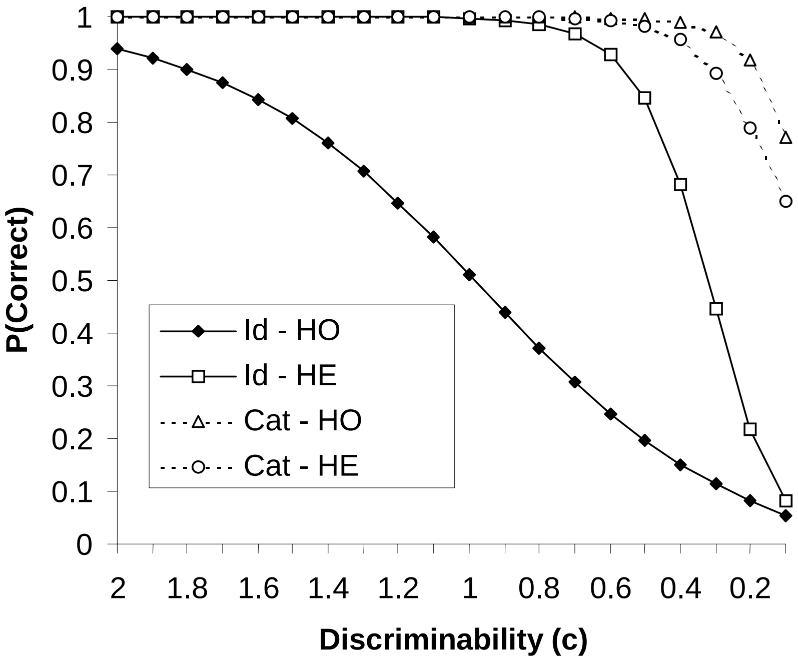
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
P(Correct)
2
1.8
1.6
1.4
1.2
1
0.8
0.6
0.4
0.2
Discriminability (c)
Id - HO
Id - HE
Cat - HO
Cat - HE

Figure 3

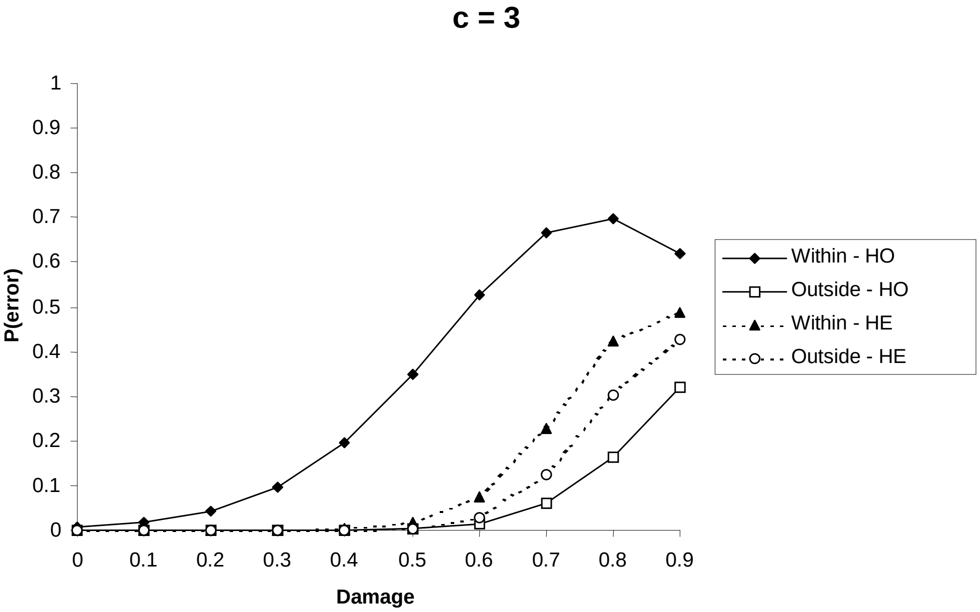
c = 3
P(error)
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
Damage
Within - HO
Outside - HO
Within - HE
Outside - HE