COGNITIVE NEUROPSYCHOLOGY, 1997, *14* (8), 1085–1130

The Interaction of Object Form and Object Meaning in the Identification Performance of a Patient with Category-specific Visual Agnosia

Mike Dixon
University of Waterloo, Ontario and the Douglas Hospital Research Centre, Montréal, Canada

Daniel N. Bub
Montréal Neurological Institute and the University of Victoria, British Columbia, Canada

Martin Arguin
Université de Montréal and Centre Hospitalier Côte-des-Neiges, Montréal, Canada

Category-specific visual agnosia following bilateral inferior temporal lobe stroke was investigated in the patient ELM. Experiment 1 verified that computer-generated blobs could not be identified when members of a set varied along a single but not along multiple shape dimensions. Experiments 2 through 6 showed that for both ELM and, to a much lesser degree, healthy participants, this dimensionality effect was modulated by semantics. By pairing the exact same shapes with semantically close vs. disparate sounds or labels, the role of an object's semantics in category-specific agnosia was assessed independently from object form. For single-dimension shape sets, the semantic proximity of the concepts associated with the shapes had no impact on ELM's identification performance. For multidimensional shape sets, ELM's error rates showed a strong positive correlation with semantic proximity ($r = .84$, $P < .01$). These results were interpreted using

Requests for reprints should be addressed to Mike J. Dixon, PhD, Department of Psychology, University of Waterloo, Waterloo, Ontario, Canada N2L 3G1 (Tel: 519 888 4567 2877; E-mail: mjdixon@watarts.uwaterloo.ca).

The authors would like to express our heartfelt appreciation to our colleague and friend ELM without whose unending patience and altruism this research could not have been conducted. The first author received support from a postdoctoral fellowship from the Alzheimer Society of Canada, and supplemental funding from the EJLB foundation. This research was also supported by grants from the Medical Research Council of Canada to the second and third authors. Martin Arguin is a Chercheur-Boursier of the Fonds de la Recherche en Santé du Québec.

an exemplar model of categorisation in which a deficit in exemplar node specificity is assumed. It is concluded that biological objects are more likely than nonbiological objects to have the combination of semantic proximity and shared values along multiple shape dimensions that pose recognition problems for patients with such specificity deficits.

INTRODUCTION

Identification deficits following brain damage may affect specific classes of objects but spare others. Patients may be able to identify accurately all manner of tools, for example, but have marked difficulty identifying even the most common animals. This tantalising phenomenon, referred to as category-specific agnosia, has generated a great deal of speculation among cognitive neuropsychologists interested in the functional architecture of object identification. If one can reliably demarcate the nature of the theoretical boundary that divides recognisable from unrecognisable objects, then we will have obtained a vital clue to some of the mechanisms underlying our ability to classify and label objects.

The observed dissociation in category-specific agnosia that has generated most research of late is one that separates biological from nonbiological categories. Although a small number of patients have been documented who show recognition deficits for man-made objects but a sparing of biological objects (Hillis & Caramazza, 1991; Sacchett & Humphreys, 1992; Warrington & McCarthy 1983, 1987), most patients with category-specific deficits show the opposite pattern; a failure to identify animal and food exemplars, but intact identification for a host of man-made artefacts. This latter deficit pattern has been documented among patients with herpes simplex encephalitis (Basso, Capitani, & Laiacona, 1988; Farah, McMullen, & Mayer, 1991; Hart & Gordon, 1992; Hillis & Caramazza, 1991; Sartori & Job, 1988; Sheridan & Humphreys, 1993; Warrington & Shallice, 1984), inferior temporal lobe strokes or closed head injuries (Arguin, Bub, & Dudek, 1996; Humphreys, Riddoch, & Quinlan, 1988; Farah, Hammond, Mehta, & Ratcliff, 1989; Etcoff, Freeman, & Cave, 1991), and Alzheimer's patients (Montanes, Goldblum, & Boller, 1995; Silveri, Daniel, Giustolisi, Gainotti, 1991; Mauri, Daum, Sartori, Riesch, & Birbaumer, 1994). It is this more prevalent pattern of deficit that is the focus of this paper.

Category-specific problems may arise from damage to different components of object recognition. Some patients may have problems accessing the stored structural knowledge of biological objects. Such patients in a reality decision task might claim that a picture of a lion with a zebra's head is a depiction of a real animal (e.g. patient Michelangelo in Sartori & Job, 1988). Others may have difficulty accessing semantic information about biological objects (e.g. patient PS in Hillis & Caramazza, 1991). Still others might have intact structural knowledge and intact semantics, yet be unable to access the correct name for biological entities in a confrontation naming task (e.g. patient TU in Farah & Wallace, 1992). Despite problems at different levels within the object recognition system, what is intriguing is that such failures appear to occur primarily for biological objects, not for man-made ones. The central purpose of this paper is to shed light on why this might be so.

Our conclusions concerning the biological vs. nonbiological distinction are based on the identification performance of the patient ELM (Arguin et al., 1996). When presented with line drawings of man-made artefacts such as tools or furniture, ELM can quickly and easily identify them. Also, when presented with the digitised sounds of animals (e.g. dog barking, lion roaring) and asked to name the animal, he is able to do so. When shown line-drawings of these same animals, however, his identification performance is extremely impaired. Thus, ELM presents with a form of category-specific identification deficit confined to visually presented objects. Consonant with his category-specific visual agnosia (CSVA) ELM fails on the reality decision test with animals, yet demonstrates intact encyclopaedic knowledge of these animals when verbally presented with their names. Thus, when relating ELM to other patients in the literature, the most appropriate comparisons are to other CSVA patients like Michelangelo, who also appear to have problems accessing the structure of objects from memory. However, on a more general level, ELM's identification performance in the present study may illuminate key differences between biological and nonbiological objects, which may have ramifications for interpreting the behaviour of all patients with category-specific deficits for living things.

Two schools of thought have emerged in attempting to explain why biological objects are harder for patients to identify than man-made objects. One school postulates that biological and nonbiological objects are each processed by their own specialised subsystems, and category-specific deficits emerge following damage to one of these subsystems. A second school maintains that a single-object recognition system processes both types of objects, but the nature of the shapes and/or the semantics of biological objects renders them preferentially susceptible to identification deficits following brain damage.

Considering first the specialised subsystem accounts, Silveri and colleagues have proposed that knowledge of biological and nonbiological categories are processed by separate systems (Silveri & Gainotti, 1988) and stored in different anatomical locations (Silveri, Daniele, Giustolisi, & Gainotti, 1991). This viewpoint would predict a sharp division between artefacts, which are recognisable, and biological objects, which are not. Contrary to this prediction, although most categories obey this theoretical boundary, there are certain man-made objects, like musical instruments, which often pose particular problems for patients who otherwise show deficits only for biological items (Damasio, 1990, Warrington & Shallice, 1984).

Sartori and Job (1988) and Sartori, Job, and Coltheart (1992) maintain that different categories of objects each have their own separate structural descriptions, and that category-specific deficits ensue following selective damage to only some of these structural descriptions. The myriad of potential semantic categories that we possess and the host of separate structural description systems required to support them prompts one to consider alternatives.

The second school of thought abandons the notion of separate subsystems for biological and nonbiological objects and views CSVA as an emergent phenomenon attributable to differences between biological and nonbiological objects in terms of their structural and semantic properties. Warrington and colleagues (Warrington & McCarthy, 1987, 1994; Warrington & Shallice, 1984) postulate that semantics is parsed into knowledge concerning sensory properties (e.g. what the object looks like) and knowledge concerning function (what the object does). If the knowledge of sensory properties becomes damaged, biological objects become unrecognisable because discrimination among exemplars relies primarily on sensory properties. Artefacts can still be recognised because patients retain knowledge of their often unique functions. Thus, category-specificity has nothing to do with whether an object is biological or nonbiological per se, but rather, whether or not an object can be defined reliably according to its function.

An alternative model is proposed by Humphreys et al. (1988), who view CSVA as emerging from an interaction of structural similarity and semantic proximity where brain damage causes recognition deficits for sets of objects whose members are both visually similar and semantically close. To account for patients primarily showing deficits for biological objects, they cite data showing that normals rate line drawings of biological objects as being visually more similar than those of artefacts. Further, in a critical experiment, judges were shown pairs of pictures and asked to rate them on a 7-point Likert scale for both their semantic and visual similarity. Based on these ratings, sets of objects varying in semantic and visual similarity were constructed. Consistent with their interaction account, the performance of the CSVA patient JB was significantly impaired on a word–picture matching test only when distracters were both visually and semantically similar to the target. When distractors were visually and semantically dissimilar, or visually dissimilar but semantically related, JB performed normally.

The proposal of Humphreys et al. provides a reasonable preliminary account for why most biological objects are unrecognisable (they form groups of structurally and semantically similar objects). In addition, these same constraints account for why certain nonbiological objects like musical instruments might also pose problems for CSVA patients (they too can form groups of objects that are structurally and semantically similar). The approach used to test the interaction hypothesis, however, suffers from a basic shortcoming that pervades almost all investigations of category-specific agnosia; the fundamental inability to specify, at the level of structure, how the actual forms of the tested objects differ from one another.

Almost without exception, researchers have used line drawings of objects to investigate CSVA. The central drawback is that the underlying shape primitives of the forms depicted by line drawings are unspecified. Although picture norms (Snodgrass & Vanderwart, 1980) are crucial for matching biological vs. nonbiological objects along factors like familiarity and image–name agreement, they are unable to specify the underlying shape primitives that allow us to distinguish between the shape of a saw and that of a cigar. Hence, exactly what mechanisms, if any, fail at the level of structural shape processing in patients who display CSVA remains a matter of conjecture.

Similarly, although it is reasonable to surmise that CSVA patients may have problems discriminating between items that are visually similar (Humphreys et al., 1988), without an understanding of a given object's shape primitives, it is difficult to know what constitutes visual similarity. That is, although ratings of similarity can easily be obtained from normals, the principles that observers use to base their judgements are unknown, hence the visual similarity that normals ascribe to a given set of objects may not be of the sort that is crucial to CSVA.

A second and arguably more important drawback in studies using line drawings is that the form of the portrayed object is inextricably yoked to the semantics of that object. Thus, if patients can identify line drawings of an axe, a pen, and a tie, but not a chicken, an eagle, and a crow, one never knows whether it is because the birds are too similar in meaning, or too similar in form. Because form is yoked to semantics, any change in one domain necessitates changes in the other domain. This inability to manipulate visual similarity and semantic proximity independently is increasingly problematic when one considers that the basis underlying normals' visual similarity ratings of line drawings is unknown.

Shape Primitives in Category-specific Visual Agnosia

The problem of line drawings having unknown shape primitives was circumvented by Arguin et al. (1996) by employing computer-generated blobs with well-defined and empirically manipulable primitives to investigate shape identification problems in the CSVA patient ELM. Arguin et al. generated shapes using three shape dimensions: Bending (B), Elongation (E), and Tapering (T). Shapes similar to those used by Arguin et al. are presented in Fig. 1. Single dimension sets were generated by assigning equally spaced values along a given dimension and generating shapes corresponding to these values. Conjunction sets, depicted in the bottom rows of Fig. 1, were generated by combining two stimulus dimensions while holding the third dimension constant. Conjunction

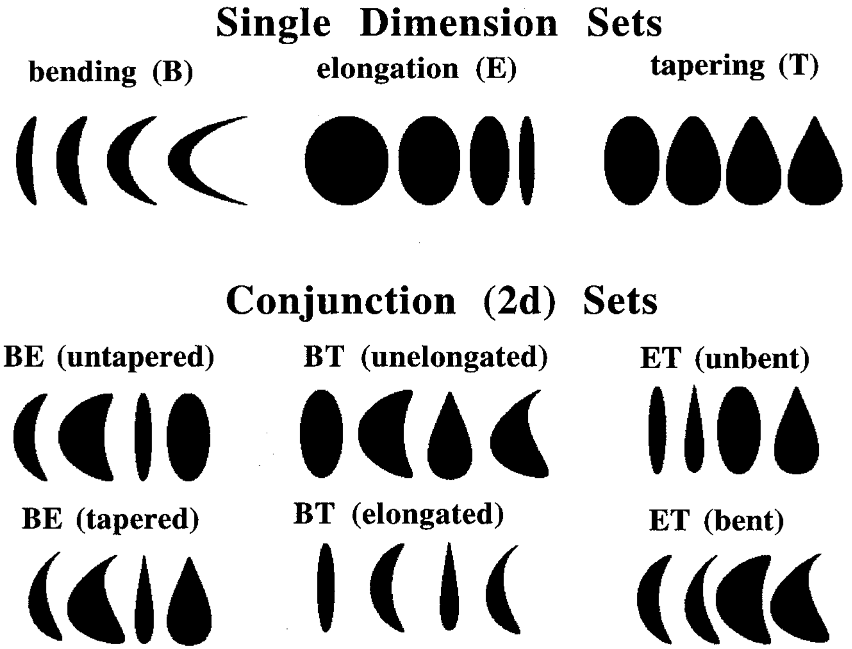

FIG. 1. Single dimension (1d) and conjunction sets (2d) used in Arguin, Bub, and Dudek, (1996) and in the current experiments 1 through 5.

sets will be referred to by the first letters of their crucial dimensions and the value of the third irrelevant dimension, hence ET bent would be shapes with the combinations of elongation and tapering, all of which are bent to an equal degree.

It will be noted that for each conjunction set there are two relevant dimensions that must be remembered in order to identify a given exemplar. Consider the BE untapered shape set in Fig. 1. In order to distinguish the banana shape from others in the set one must remember that it is both bent and elongated. If one remembers only bending, then the two bent shapes will get confused in memory; if one remembers only elongation then the two elongated shapes will get confused. Thus, in order to disambiguate the members of this set, values on two shape dimensions must be retrieved because for every shape there is another shape in the set that has the same value on one of these critical dimensions. In the single dimension sets, however, there is only one relevant dimension, and there are no shared values along this relevant dimension.

In order to investigate whether shapes that shared multiple visual dimensions posed particular problems for ELM, Arguin et al. devised the following procedure. Four blobs comprising either the single dimension or conjunction set were simultaneously presented in the four corners of the screen for a limited duration. On test trials one of these blobs was centrally presented and ELM was asked to point to this blob's former location. ELM was consistently better able to perform this pointing task with items from single dimension sets (29.2% errors) than when given items from conjunction sets (56.7% errors). Importantly, Arguin et al. showed that in a perceptual matching-to-sample task where memory requirements were minimised, ELM made more errors on the more perceptually similar single dimension sets (8.8%) than on the conjunction sets (0.004%) errors.

These and related experiments indicated that ELM's problems with shape identification involve memory for the visual properties of objects, not their perception. Specifically, he has difficulty extracting from memory information about *multiple* critical shape dimensions (e.g. as a real-world example, he knows a banana is long and thin, but cannot remember if it is tapered or curved). Although certain dimensions appeared easier for ELM to remember than others, in general it did not matter which shape dimensions were combined to form a conjunction set; as long as values along more than one dimension had to be extracted from memory to disambiguate exemplars, ELM's performance suffered dramatically.

Importantly, ELM's deficit was shown to interact with object semantics. In a reaction time experiment Arguin et al. required ELM to name single dimension and conjunction sets of blobs as quickly as possible. In one condition identifications were made using fruit and vegetable labels corresponding to what the blobs looked like. For example, the blobs in the ET unbent set in Fig. 1 were given the labels “cucumber”, “carrot”, “melon”, and “pear”. In a second condition the exact same blobs were given artefacts labels that also bore a resemblance to the blob's form (cigar, tent-pen, balloon, and spinning-top). For fruit and vegetable labels, ELM's reaction times followed the expected pattern; single dimension sets were named significantly faster than conjunction sets. For the same blobs named using artefact labels, however, conjunction set reaction times were as fast as those for single dimension sets, and significantly faster for the same shapes labeled using fruit and vegetable names. Thus, these results provided preliminary evidence that ELM's difficulties with sets of objects sharing values along critical shape dimensions may depend upon the semantic proximity of the labels applied to the shapes being identified.

ELM and Models of Category Learning

Although the behaviour of ELM in the Arguin et al. study is difficult to reconcile with any of the extant neuropsychological theories of CSVA, it can be readily understood within the framework of some exemplar models of category learning in normals (Estes, 1994; Kruschke, 1992; Nosofsky, 1986). These models were developed to explain the manner in which normals categorise objects, rather than identify them. Technically, however, identification and categorisation differ only in a very elementary statistical sense: for identification every exemplar is assigned a unique response, whereas for categorisation many examplars can be mapped to the same response. Thus, highly similar exemplars will often lead to correct categorisation, but for identification the greater the exemplar similarity the greater the probability of identification errors.

In Kruschke's (1992) ALCOVE model, there are three layers: input nodes, hidden exemplar nodes, and output nodes. Each input node encodes stimulus values on a single psychological dimension (e.g. bending, elongation, and tapering in Fig. 1). The hidden exemplars nodes, to which inputs are connected, are represented as points in a multidimensional psychological space. The location of these hidden nodes corresponds to the various combinations of shape dimension values that comprise the exemplars within a set. For example, in an ALCOVE model designed to identify the ET unbent set in Fig. 1 there would be hidden nodes located at [0,0,0], [0,0,1], [0,1,0], and [0,1,1] within a psychological space in which the location coordinates are values on bending, elongation, and tapering, respectively. In this multidimensional psychological space the objects [0,0,0] and [0,0,1] would lie closer to one another (differing only in tapering), than the objects [0,0,0] and [0,1,1], which differ on both elongation and tapering. Hidden exemplar nodes have activation profiles. They respond most strongly to stimuli which have the same values as their location coordinates (e.g. the exemplar node located at [0, 1, 1] would respond maximally to a presented shape comprised of bending = 0, elongation = 1, and tapering = 1, but still would respond strongly to exemplars comprised of similar values [e.g. 0, .9, 1]). Activation falls off exponentially as similarity between the input stimulus and the exemplar node location decreases. The spatial extent of these activation profiles (i.e. the "receptive field") of these hidden exemplar nodes depends on a specificity parameter. Large specificities mean that hidden nodes will respond only to stimuli very close to the exemplars that they code. Thus, the [0, 1, 1] shape would respond to shape [0, .9, 1.1], but not to shape [0, .5, 1]. Small specificities mean large receptive fields. Thus, the exemplar node located at position [0, 1, 1] would become activated by both shape [0, .5, 1] *and* shape [0, 0, 1] but perhaps not [0, 0, 0], which is located too far away in psychological space. Output nodes, representing responses, have learned connection strengths to the exemplar nodes.

ALCOVE can account for a wide range of categorisation phenomena because the connections between the input nodes and the exemplar nodes are gated by an attentional dimension strength. This gating mechanisms acts as a multiplier, capable of increasing or decreasing the strength with which a given dimension connects to exemplar nodes. Over the course of learning, performance is optimised by placing greater weights on those stimulus dimensions that best define categories (for identification each exemplar is a separate category). In ALCOVE, learning to identify the ET unbent shapes would be optimised by decreasing the attention strength on the uninformative bending dimension, but increasing strengths on the two relevant dimensions. Similarly for a single dimension set, performance would be optimised by increasing the strengths for the single relevant dimension and decreasing the weights for all other dimensions. ALCOVE mimics the performance of normals in that categorisations based on single dimensions are learned more quickly than categorisations based on multiple dimensions. In ALCOVE the greater the number of relevant dimensions, the more difficult the categorisation (or identification) problem.

ELM's blob identification performance can be interpreted most parsimoniously by assuming a deficit in the specificity parameter. His deficit is that he has much larger exemplar-node receptive fields than normal. These wide receptive fields tend to overlap for objects that are close together in psychological space, and these overlapping receptive fields cause ELM to confuse these closely located objects. In our ET unbent conjunction set example, the hidden exemplar node that responds maximally to the [0,1,1] shape would still become activated by [0, 0, 1] shape, causing ELM to confuse these items when trying to identify them. Although wider than normal, the patterns of his confusions indicate that his exemplar-node receptive fields are not infinitely wide. That is, he seldom confuses the [0,1,1] shape with the [0,0,0] shape.

ELM can compensate for wide activation profiles when exemplars can be distinguished using a single dimension by increasing the attentional dimension strength for that dimension and gating the other irrelevant dimensions. However, when objects within a set share values along multiple dimensions (as in conjunction sets), this gating strategy is not effective and his performance suffers dramatically[1].

---

[1] A simple computer simulation for identifying single dimension and conjunction sets of blobs was conducted using the ALCOVE architecture. The single dimension set include four equally elongated shapes varying in bending. Two stimulus dimension nodes, four hidden exemplar nodes, and four response nodes were used. To mimic normal single dimension set performance, attentional dimension strengths were set by hand at 4.0 for bending and 0 for elongation, to reflect performance optimisation by attending to the single relevant dimension of bending. The specificity parameter that controls the hidden exemplar-nodes receptive field size was set by hand at $c = 8$ (high specificity–small receptive fields). Stimuli varying on this dimension were presented to this architecture. Results of this simulation were 99% correct detections. To simulate ELM's impairment, the specificity parameter was set to ¼ that of normal ($c = 2$). This resulted in 81% performance (approximately the single dimension set performance ELM showed in Arguin et al., 1996).

For conjunction sets attentional dimension strengths were set to 2.0 for each stimulus dimension. The model was presented with shapes from the BE untapered set. Results of the simulation revealed 97% correct identifications when specificity was high ($c = 8$). Once again to mimic ELM's impairment, phi was reduced by ¼. With $c = 2$, the model correctly identifies a given exemplar 56% of the time. These results are consonant with the error rates observed by ELM in Arguin et al. (1996).

Although such a model can account for ELM's single dimension and conjunction set performance in tasks devoid of semantic content, it must be recalled that semantics can modulate his conjunction set performance (having ELM label the ET unbent blobs, a cigar, a tent-peg, a balloon, and a spinning top leads to fast, nearly error-free performance).

In order to account for semantic modulation of conjunction set errors, all that is required is the reasonable assumption that the label borne by a shape is a contributing determinant of where a hidden exemplar node is located in psychological space. Thus, in this new conception of ALCOVE, the location of the hidden units are determined not only by structural information (Bending, Elongation, Tapering), but also by semantic information. One can assume that, within a multidimensional space that uses both structural and semantic attributes as coordinates for exemplar locations, semantically related objects are located close to one another whereas semantically unrelated objects are located further apart. Thus, in this psychological space, the hidden exemplar nodes for objects labeled as carrot and cucumber would lie closer together than the exact same blobs labeled tent-pen and cigar. Since closer objects are more confusing, the former should take longer to disambiguate than the latter.

The following series of studies, conducted with the help of the CSVA patient ELM and guided by exemplar models like ALCOVE, will look at the influence of both structural and semantic factors on exemplar identification. In so doing we will address both the problem of line drawings having unknown underlying primitives and the problem of semantics being yoked to object form. The first problem, concerning the unspecified nature of line drawings, will be addressed by using Arguin et al.'s computer-generated shapes. In Experiment 1 these shapes will be used to replicate Arguin et al., and provide evidence that the CSVA patient ELM has problems disambiguating exemplars from sets of objects that share values on multiple critical shape dimensions. In Experiments 2 through 6 these same shapes will then be used to decouple the influence of semantics in object recognition from the influence of object form. Shapes will arbitrarily be associated with concepts that are either semantically close (e.g. robin, crow, owl, turkey) or semantically disparate (helicopter, telephone, saw, tennis-racquet). Based on the notion that ELM has problems disambiguating exemplars that lie close to one another in a multidimensional psychological space (because of abnormally wide exemplar-node receptive fields) it is predicted that identification performance will suffer the most when objects within a set have both overlapping visual features *and* refer to semantically close concepts. By using the same computer-generated shapes to stand for both semantically close and semantically disparate concepts we will present a paradigm which, for the first time in Neuropsychology, allows the structure of objects to be held constant, and the effects of semantic proximity on object recognition to be directly evaluated.

EXPERIMENT 1

The Nature of the Structural Processing Deficit in ELM

The exemplar model of category-specific visual agnosia proposed earlier is somewhat similar to that of Humphreys et al.'s (1988) in that both theories reject the notion of specialised subsystems for biological and nonbiological categories of objects, interpreting category specificity instead as an emergent phenomenon based on the structural and semantic properties of the objects being recognised. The primary difference between the two accounts concerns the notion of visual similarity. For Humphreys et al., visual similarity is operationalised by normals' ratings of how similar objects appear. In the exemplar model account described earlier, visual similarity depends on the number of shape dimensions that must be extracted from memory in order to identify a given member of a set uniquely.

The single dimension and conjunction sets used by Arguin et al. may prove heuristically useful in empirically distinguishing between these two accounts of CSVA. Humphreys et al. surmised that sets of objects which normals judge to be more structurally similar would pose greater recognition problems for CSVA patients than less structurally similar sets. Looking at Fig. 1, however, the single dimension sets appear to be more structurally similar than the conjunction sets, but as Arguin et al. have repeatedly shown, ELM has a much easier time distinguishing exemplars from this set than from the more structurally distinct conjunction sets.

Experiment 1 will serve as a replication of Arguin et al., using a slightly different paradigm. Instead of presenting all four exemplars on the screen at once, on learning trials, exemplars will be presented one at a time. (This will prevent ELM from making simultaneous comparisons of all set members. Instead he will have to encode set members sequentially—a situation that is arguably more analogous to veridical object learning.) In addition, Humphreys' notion of structural similarity will be pitted against our exemplar model. It is predicted that normal raters will rate a single dimension set as being more visually similar than a conjunction set but, in contrast to what must be predicted by Humphreys et al., ELM will demonstrate fewer identification errors for this single dimension set than for a more structurally distinct conjunction set because he can compensate for his abnormally wide receptive fields by gating irrelevant dimensions and increasing the attentional dimension strength on the single relevant dimension.

Method

ELM

Clinical history. ELM was born in 1928. Now retired, he formerly worked in the purchasing department of a manufacturing plant. In December of 1982 ELM was admitted to hospital for heart failure. Neurological symptoms of sudden onset were reported on 5 December 1982. These included: nominal dysphasia, left/right confusion, dyscalculia, and agraphia without alexia. An emergency CT scan revealed a hypodensity deep in the right mesiotemporal lobe. The neurological symptoms resolved and upon discharge ELM suffered from a residual nominal aphasia and mild memory impairment, which later disappeared. In August of 1985 he was readmitted to the Montreal Neurological Hospital, presenting with pronounced anomia, memory impairment, and dysgraphia. A CT scan conducted on 9 August revealed irregular enhancing lesions deep in the left and right mesiotemporal lobes. His condition improved and he was discharged on 21 August 1985.

Neuropsychological assessment. In October of 1987 ELM underwent neuropsychological testing, which revealed normal IQ (93 WAIS-R verbal, 91 WAIS-R performance) but residual impairments in the delayed recall of both verbal (WMS verbal = 10.5) and pictorial material (WMS recall of geometric forms = 1). Also, he showed impairment in visual object recognition (Wingfield Object Naming 11/26) and face recognition (Benton Facial Recognition Task = 33). In clinical testing his object recognition deficit seemed to be attributable to an impairment in identifying pictures of animals.

A more in-depth analysis of his visual recognition deficit using Snodgrass and Vanderwart (1980) pictures revealed a marked discrepancy between biological and nonbiological objects in confrontation picture naming. He correctly identified only 21% of biological items, but correctly named 92% of man-made artefacts. Although name frequency and familiarity were both significant predictors of his naming performance, the biological–nonbiological distinction was the strongest predictor of naming accuracy within a multiple regression framework.

In a reality decision task ELM was shown pictures of stimuli and asked whether each one was real or not. Negative items were created by interchanging parts (a cow's body with a dog's head). ELM could make reality decisions with objects (38/41) but not with animals (41/70).

Although ELM's ability to recognise pictures of objects is impaired, his encyclopaedic knowledge of them is intact. For example, when given the word *camel* and asked to define what it was, ELM said "it is an animal that more or less lives in the Sahara desert. Some people refer to it as 'the ship of the desert.' It can go for days without drinking water."

Importantly, ELM's perception is intact. He can copy both complex geometric forms (Copy portion of Rey's Complex Figure = 31/33) and animals. He is normal in naming photographs of household objects taken from both canonical (26/27) and noncanonical views (25/27). He can also match canonical and noncanonical views of animals (7/7) and artefacts (18/19). He shows normal global to local interference for Navon Stimuli, and has no problem identifying overlapping objects.

Control Participant(s)

A female, age-matched control was used in the shape location phase of this experiment. In addition, five university students were asked to provide pairwise similarity ratings for the single dimension and conjunction sets used in the experiment.

Materials and Procedure

Stimuli. Computer-generated shapes with specified values of bending, elongation, and tapering were used throughout this experiment. Stimuli consisted of black blobs displayed against a white background. The conjunction set comprised four items, which varied in elongation and bending but had consistent values on tapering. The single dimension set comprised four items, which varied only in elongation. All stimuli were presented using a Macintosh Quadra computer interfaced to a AppleColour high resolution RGB monitor.

Similarity ratings. Five independent judges were each shown the eight shapes comprising the single dimension (four shapes) and conjunction (four shapes) sets. Shapes were shown one at a time, with the four shapes comprising the single dimension sets displayed first, followed by the four shapes comprising the conjunction set. Each set was shown three times. Following this familiarisation phase, participants were asked to give "visual similarity" ratings to pairs of shapes using a 7-point Likert scale ranging from 1 (very dissimilar) to 7 (very similar). Twelve pairs of shapes (six pairs comprising the single dimension set, and six pairs comprising the conjunction set) were rated. These 12 pairs were presented in random order.

Shape–Location Learning

Learning trials. Prior to the experiment, each of the four shapes within a given set were randomly assigned a location on the computer screen (top, bottom, left, or right). These shape–location pairings remained constant throughout the experiment. To prevent participants from using local environmental cues (e.g. tapered shapes always pointing in a given direction), on a given trial the shape would appear in its assigned location using one of eight orientations (rotations of 0, 45, 90, 135, 180, 225, 270, 315 degrees from vertical). On learning trials, shapes were presented in their respective locations one at a time (see Fig. 2). Shapes remained on the screen until the participant indicated they had adequate time to view the stimulus and memorise its location. Eight learning trials (two of each shape) were presented, following by eight test trials (two of each shape).

Test trials. On test trials, shapes were presented in the middle of the screen (using one of the eight orientations) and ELM was asked to name its former (learning-trial) location (see Fig. 2). This pattern of 8 learning followed by 8 test trials was repeated 12 times, for a total of 96 learning and 96 test trials. After a 5-minute break a second, identical block of 96 learning and 96 test trials was administered. Shapes were presented at each of the eight orientations an equal number of times.

The single dimension and conjunction sets were run on separate days.

Results

Similarity Ratings

For each shape set, there are six pairwise contrasts to be made among the four shapes (AB AC AD BC BD CD). Thus, each participant made six pairwise

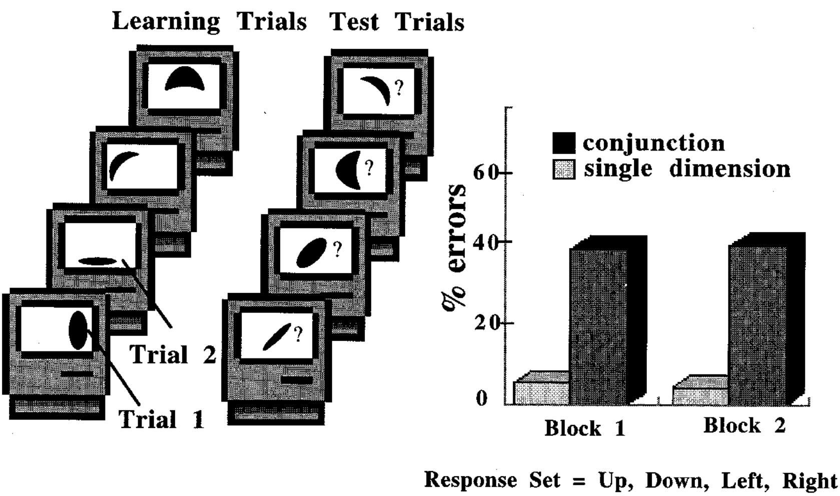

FIG. 2. Shape-location paradigm and Experiment 1 results. Shapes were displayed in a fixed, preassigned location on learning trials. On test trials the blob was presented in the centre of the screen and the subject attempted to name the location associated with that blob. Results for block 1 and 2 are presented for single dimension and conjunction sets.

ratings for single dimension sets and six pairwise ratings for conjunction sets. The data of the 5 participants were pooled to yield 30 (6×5 participants) pairwise ratings of the single dimension set and 30 pairwise ratings of the conjunction set. Participants gave significantly higher similarity ratings to the single dimension set (mean = 3.6) than the conjunction set (mean = 2.4) [$t(29) = 2.17$, $P < .05$].

Shape Identification

In assessing the significance of shape identification trials through all experiments, $\alpha = .01$.

Control. For both the single dimension and conjunction sets the age-matched control participant had correctly paired the shapes with their locations after the first set of 8 learning trials (0 errors on 96 trials). Because of perfect performance on block 1 the control was not tested on block 2. Because of the absence of errors the control data will not be compared statistically to ELM's performance but will serve only to indicate how easy the task is for normal participants.

ELM. For shapes varying along a single dimension, ELM made 5/96 (5.2%) errors on block 1 and 4/96 (4.2%) errors on block 2. On the conjunction set, ELM made 39/96 (40.6%) errors on block 1 and 40/96 (41.6%) errors on block 2. These results are depicted in the right-hand panel of Fig. 2. ELM's performance was significantly poorer for conjunction sets relative to single dimension sets for both block 1 ($\chi^2 = 26.27$, $P < .001$), and for block 2 ($\chi^2 = 29.45$, $P < .001$) respectively.

Discussion

The Nature of the Structural Processing Deficit in ELM

The markedly poorer performance of ELM on the conjunction set, relative to the single dimension set[2], indicates severe problems remembering shape–

[2]The single dimension set error rate is somewhat lower in this location experiment than in Arguin et al. (1996). In Arguin et al., single dimension sets contained a blob that had a zero value on the relevant single dimension (using the example for the elongation set they used a blob that was circular, along with three blobs of increasing elongation; for bending, an unbent blob along with three blobs of increasing bending). Such a situation may have caused ELM to treat these one-dimensional sets as having more than one critical dimension. In the bending set, ELM may have coded the unbent shape as being elongated, and the other three as bent. In the present experiment, blobs with zero values were avoided, a situation that may focus attention unequivocally on a single dimension.

location pairings when identification of a given shape involves extracting from memory specific values along more than one relevant shape dimension.

It should be noted that ELM's problem is not one involving the ability to perceive the differences between these shapes. In a previous experiment by Arguin et al., ELM actually showed better performance for conjunction sets relative to single dimension sets when the memory components of the task were minimised. Hence ELM's problem is not one of perception per se, but rather reflects a post-iconic memory problem for conjunctions of shape features.

Humphreys et al. (1988) argued that the specific types of stimuli that would pose problems for CSVA patients are objects that are structurally and semantically similar. Although this particular experiment does not deal with the semantic component of this interaction hypothesis, it has a great deal to say about structural similarity. Using the same 7-point Likert scale rating procedure as that employed by Humphreys et al., the set of shapes that normals rated as being visually *more* similar posed *less* of an identification problem for ELM than the set whose members were rated as being relatively dissimilar.

Such a finding suggests that the visual similarity judgements of normals do not always accurately reflect the underlying principles that determine shape identification problems in a category-specific agnosia patient like ELM. An exemplar model that assumes abnormally wide receptive fields for ELM can, however, account for why he has trouble with conjunction sets—exemplars within the set are stored close together in multidimensional space, and the abnormally wide receptive fields for these closely stored exemplars tend to overlap—a situation that generates confusions among these closely stored objects. More importantly, we can explain why objects within a single dimension set that was rated as visually closer than the conjunction set posed fewer identification problems: ELM can disambiguate exemplars by gating irrelevant dimensions and increasing the attentional dimension strength on the single relevant dimension.

The relevance of these distinctions to everyday object recognition can be made clear by considering the set: cup, bowl, glass, and vase. This constitutes a visually similar set of objects that is also semantically similar. As such, Humphreys et al. must predict that a CSVA patient like ELM would confuse these objects. Our exemplar model, by including dimensional gating mechanisms, explains why ELM is able to identify these objects flawlessly in real life; they vary only on the single dimension of elongation.

From Structure to Meaning: The Role of Semantic Proximity in CSVA

Arguin et al. (1996), and Experiment 1, demonstrated that ELM has problems in identifying shape sets in which exemplars shared values on critical shape dimensions. As previously mentioned, this cannot in and of itself account for category-specific visual agnosia. After all, the shapes of a balloon, cigar, spinning top, and tent-peg form a conjunction set (ET unbent) where exemplars share values in critical shape dimensions, yet ELM can quickly and correctly label these shapes. Semantics therefore, *must* play a critical role in CSVA.

Although at first it may appear that Experiment 1 is devoid of semantic content, it must be noted that in this paradigm the category labels (up, down, left, right) are semantically related (all are directions, locations on a computer screen, etc.) These semantically related labels may serve to heighten the similarity among the four exemplars, relative to a paradigm in which shapes were associated with completely unrelated labels (e.g. "cigar", "spinning top", "tent-peg", and "balloon" used by Arguin et al., 1996). In all ALCOVE-type models, in which the locations of hidden exemplar nodes within a multidimensional space depend upon both visual information and semantic information, the conjunction set of blobs whose semantics involve four computer screen locations would be stored closer together in this space than a conjunction set of blobs whose semantics are unrelated (cigar, tent-peg, etc.) Because closely stored objects are more confusing than objects stored further apart, for those like ELM, with abnormally wide exemplar-node receptive fields, it follows that they would have a much more difficult time with the location task than a blob-labeling task in which labels are semantically unrelated.

If these contentions are correct, then a number of hypotheses can be generated. First, because of intact attentional gating mechanisms, ELM should be able to disambiguate the exemplars of single-dimension sets regardless of their semantics. Second, when sets of objects are used that preclude the use of attentional gating (e.g. conjunction sets), ELM's performance should depend on the semantics associated with the exemplars within a set. For sets of objects that have overlapping visual features *and* overlapping semantics, ELM should show massive confusions. For similar shapes having disparate semantics, however, fewer identification problems should ensue because the disparate semantics serve to provide greater separation of these objects within multidimensional psychological space.

In order to test these hypotheses one could search for semantically proximate and semantically disparate concepts whose shapes resemble the single dimension and conjunction sets in Fig. 1. Such an approach is ill advised because few of the blobs that can be generated by varying elongation, bending, and tapering actually look like real-world objects. An alternative approach is to abandon any real-world correspondence between the blob's form and the label we asked ELM to ascribe to it. This can be done by repeatedly accompanying *blob A* with the sound of a dog barking; *blob B* with a wolf howling, *blob C* with an elephant trumpeting, etc. Thus when ELM labels these shapes he will simply name the sound associated with that shape. The advantage of not having any correspondence between the shape of the blob and the label associated with it is that the exact same set of blobs can be paired with either semantically similar (e.g.

robin's song, crow's caw), or semantically disparate experiences (e.g. sound of a helicopter, sound of a saw, etc.) Thus, form can essentially be held constant and the role of object semantics in CSVA assessed independently.

EXPERIMENT 2

Method

Subjects

ELM and the same age-matched control in Experiment 1 were retested in this experiment.

Materials

Stimuli. Similar computer-generated shapes comprising single dimension and conjunction sets as described in Arguin et al. (1996) and in Experiment 1 were employed in this experiment.

Shape sets and sound pairings. The following sets of sound-shape pairings were used:

1. The semantically disparate sound of a leaf-blower[3], a telephone, and water pouring into a glass, were paired with the conjunction set (BE untapered) and the single dimension set B.
2. The semantically close sounds of a dog barking, a horse neighing, a wolf howling, and an elephant trumpeting were paired with the conjunction set (BE tapered) and the single dimension set (E).
3. The semantically disparate sounds of a saw, a tennis-racquet, hitting a tennis-ball, a photocopier, and a helicopter were paired with the conjunction set (ET not bent), and the single dimension set (E).
4. The semantically close sounds of a robin, a crow, an owl, and a turkey were paired with the conjunction set (ET unbent) and the single dimension set (B).
5. The semantically disparate sounds of a saw, a tennis-racquet, a photocopier, and a helicopter were paired with the SET 4 shapes (ET, unbent).

Each shape set is depicted in Fig. 1.

[3]The original sound recording was a somewhat distorted recording of a wasp. On sound identification trials ELM insisted that this recording sounded more like a leaf-blower than a wasp, so we used this concept instead.

Sound identification trials. Prior to testing a given shape set, ELM was presented with the digitised sound recordings unaccompanied by shapes. Each sound recording was presented 6 times for a total of 24 trials and ELM was asked to match the sound with one of sound names written on an index card.

1. *Learning trials:* On a given trial a single shape would appear in the centre of the screen using one of eight orientations (rotations of 0, 45, 90, 135, 180, 225, 270, 315 degrees from vertical). Along with the shapes a cursor (in the form of a pointing finger) would appear. ELM was instructed to place the cursor over the shape and click the mouse, whereupon the sound recording would be played over two small speakers. Shapes remained on the screen during the playing of the sound recordings and stayed on until ELM indicated he had had time to memorise the shape–sound pairing. Eight learning trials were presented (two of each shape) in random order with the restriction that identical shapes did not immediately follow one another.
2. *Test trials:* Following the eight learning trials, eight test trials were presented. On test trials, shapes were presented in the middle of the screen (using one of the eight orientations) and ELM was asked to name the sound associated with it (e.g. saw, robin, etc.) An index card containing the names (e.g. "ROBIN", "CROW", "TURKEY", "OWL") of the four sound alternatives was provided for ELM to refer to as needed. The 8 learning–8 test trial pattern was repeated 12 times, for a total of 96 learning and 96 test trials. After a 5-minute break a second identical block of 96 learning and 96 test trials was administered. Shapes were presented at each of the eight orientations an equal number of times. The procedure for learning and test trials is depicted in the upper left panel of Fig. 3.

To avoid interference effects caused by the same shapes standing for more than one concept, on any given day only two shape sets were administered. These sets used different shapes and different sound assignments. Conjunction sets and single dimension sets were run on alternate days. The order of set presentation was counterbalanced. On day one the (semantically disparate) conjunction set BE untapered was run first followed by the (semantically close) BE tapered set. On day two the (semantically close) single dimension set E was run followed by the (semantically disparate) set B. On day three conjunction sets were run in the reverse order relative to day one: (semantically close) ET unbent was run first followed by the (semantically disparate) ET bent set. On day four the single dimension sets were run in reverse order relative to day two: (semantically close) set B, followed by (semantically disparate) set E.

Finally, a third, semantically disparate, set was run several weeks later (saw, tennis-racquet, photocopier, helicopter) using the exact same shapes (ET unbent) that were used in set 4 (which was previously paired with the four bird songs).

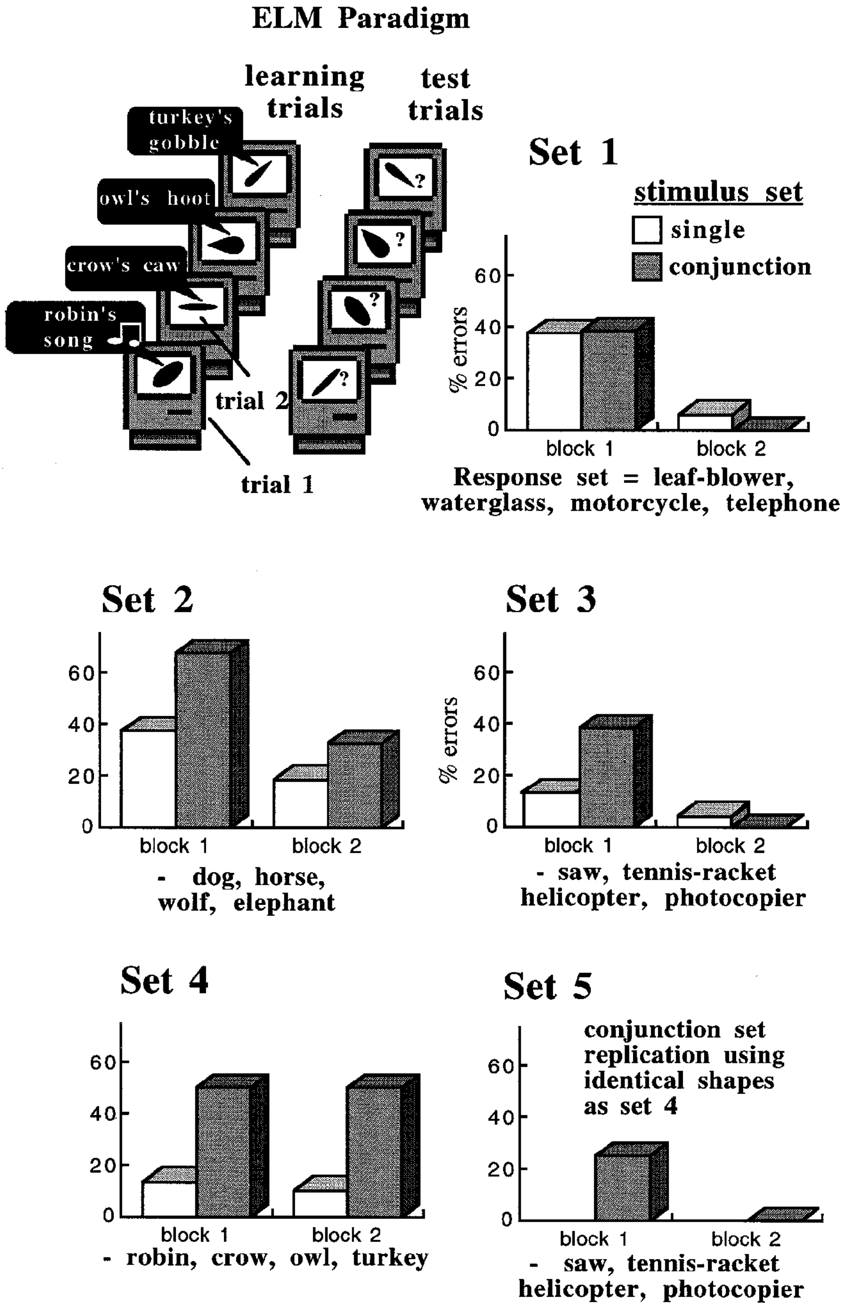

FIG. 3. ELM's block 1 and 2, error percentages for single dimension and conjunction sets when shapes were paired with semantically close or disparate concepts (Experiment 2)..

Results

ELM showed flawless performance in the sound identification trials (24/24 correct identifications on each of the four sets of sounds.

Shape Identification

Control. Error rates for the age-matched control were once again at or near floor levels. Of the few errors that were made all occurred during the first 16 test trials. The remaining 80 test trials were error free. For this reason only 1 block of 96 learning and 96 test trials were run for the control participant for all sets. For conjunction set 1 (semantically disparate) 3/96 errors were made. On conjunction set 2 (four semantically close animals) a single error was made, and the two remaining conjunction sets were error free. For single dimension sets performance on set 1 (semantically disparate) was 6/96 errors, on set 3 (4 semantically close birds) a single error was made, and on sets 2 and 4 performance was flawless. Once again, because of the restricted range of errors, her data were not analysed, but served only to indicate the ease with which a non-brain-damaged individual can perform these blob-labeling tasks.

ELM. Error rates for ELM were as follows:

1. *Conjunction sets*

a. *(Semantically close):* Error rates ranged from 50% to 68% to on block 1 and 32% and 50% on block 2. These results are depicted by the dark bars in sets 2 and 4 of Fig. 3. Combining results from these semantically close sets, error rates were (117/192 = 60.94%) on block 1 and (79/192 = 41.14%) on block 2.

b. *(Semantically disparate):* Error rates ranged from 30.2% to 38.5% on block 1, to 0% (both sets) on block 2. These results are depicted by the dark bars in sets 1 and 3 of Fig. 3. Combining there sets, error rates were (66/192) = 34.37% on block 1 and (0/192) = 0.0% on block 2. Using the combined data sets and comparing semantically close to disparate sets, the conjunction sets showed significantly higher error rates for semantically close sets on both block 1 (60.94% vs. 34.37%, $\chi^2 = 7.39$, $P < .01$), and block 2 (41.14% and 0% $\chi^2 = 41.14$, $P < .0001$).

2. *Single dimension sets*

a. *(Semantically close):* Error rates ranged from 10.4% to 37.5% on block 1 and between 10.4% and 18.8% on block 2. These error rates are depicted by the light bars in sets 2 and 4 of Fig. 3. Combining these semantically similar single dimension sets, error rates were 25.56% on block 1 and 14.58% on block 2.

b. *(Semantically disparate):* Error rates ranged from 16% to 37.5% on block 1, and from 1% and 6% on block 2. These error rates are depicted by the light bars in sets 1 and 3 of Fig. 3. Combining these semantically disparate single dimension sets errors were 26.56% on block 1 and 3.64% on block 2. Using the combined data sets and comparing semantically close sets to semantically disparate sets, no significant differences were obtained for single dimension sets. Error percentages were not significantly different between semantically disparate and close sets on block 1 (25.52% vs. 26.56% on block 1, $\chi^2 = 0.02$, n.s.) or between close and disparate sets on block 2 (14.58% and 3.64%, $\chi^2 = 6.567$, n.s.)

3. *Identical shape comparison.* Finally, the two conjunction sets (semantically close and disparate) involving the exact same shapes were compared (sets 4 and 5). For the semantically similar set 50% errors were recorded for both blocks 1 and 2. For the semantically disparate set errors were 25% and 0% for blocks 1 and 2 respectively. Significant differences between semantically close and disparate sets were noted for both first block ($\chi^2 = 8.0$, $P < .01$) and second block performance ($\chi^2 = 48.0$, $P < .0001$). ELM's semantically close and disparate conjunction set naming performance for these identical shape sets is depicted in sets 4 and 5 of Fig. 3.

Discussion

For the normal participant this task was considered trivial. Only once did she require more than two eight-trial learning blocks to successfully map the four sounds to the four shapes. Typically she had mapped shapes to sounds after the first eight learning exposures, resulting in error-free performance over 96 trials. In exemplar model terms, the ability to disambiguate individual exemplars from sets of objects having overlapping shape dimensions and semantic attributes reflects highly specific, narrow, exemplar-node receptive fields. That is, even objects that are stored close together in multidimensional space do not typically generate confusions, because narrow receptive fields associated with these exemplars do not overlap in healthy observers.

For ELM, who we presume has abnormally wide receptive fields, Experiment 2 provides empirical evidence that in category-specific visual agnosia, identification difficulties are modulated by both visual shape dimensions *and* semantic proximity. For single dimension sets where there are no shared values among exemplars, the effect of semantic proximity is minimised because ELM can increase the attentional dimension strength of the single relevant dimension and disambiguate exemplars at reasonable levels of proficiency even if these objects are semantically quite close.

For conjunction sets, however, this gating strategy is ineffective because multiple dimensions are required for exemplar disambiguation. Thus, conjunction sets create the *potential* for recognition difficulties. If objects are psychologically close by virtue of having both shared visual features *and* semantically overlapping attributes then profound recognition problems ensue. When, however, such exemplars are associated with experiences that are semantically disparate, this serves to separate these objects within psychological space enabling performance (0% errors on blob 2) to climb to single dimension set levels (3.64% errors).

The research of Humphreys et al. (1988) has shown that interactions involving structural similarity and word frequency can occur. Although no attempt was made to match disparate and close sets expressly in terms of word frequency, this variable does not seem to bear any relation to ELM's error rates. For semantically close sets, ELM showed poor performance for both high-frequency (dog, horse, wolf, elephant: mean frequency = 94) and low-frequency sets (robin, crow, owl, turkey: mean = 3.25). As such the effects of word frequency, if any, were completely subjugated by the overwhelming effects of semantic proximity.

Further, it is important to note that the pattern of ELM's errors cannot be attributable to practice or fatigue effects, as the order of semantically close and disparate pairings was counterbalanced. In addition, the semantic effect on conjunction sets cannot be attributable to differences between the conjunction sets employed; sets differed only by zero or non-zero values on an irrelevant dimension (e.g. BE untapered or BE tapered) and conjunction sets of both kinds were each associated with perfect block 2 performance in the semantically disparate condition. Conclusively, good performance for conjunction sets associated with disparate concepts, and poor performance for sets associated with semantically close concepts, was noted using *exactly the same shape sets*. This latter finding provides the strongest support for a deleterious effect of semantic proximity on object identification performance in ELM. Holding shape constant in this fashion is especially important when one considers that Arguin et al. (1996) found that ELM was preferentially sensitive to certain dimensions (e.g. changes in elongation appeared to be more salient to ELM than changes in curvature or tapering).

The error performance of ELM on this blob-identification task replicates the reaction time performance of ELM in Experiment 6 of Arguin et al. (1996). In the Arguin et al. study, blobs were given fruit and vegetable labels corresponding to the blob's form (cucumber, carrot, melon, and pear for the ET unbent set in Fig. 1). In a second condition the same blobs were given artefact labels that also bore a resemblance to the blob's form (cigar, tent-peg, balloon, and spinning-top). For the conjunction set mapped to semantically disparate labels, ELM was significantly faster at identifying these blobs than when semantically close labels were used. The reason Arguin et al. used reaction time rather than identification performance (as in the current study) was that, unexpectedly, ELM made very few identification errors even in the fruit and vegetable condition. Why would ELM be almost error free on the ET unbent set using fruit and vegetable labels, yet have such difficulty with exactly the same shapes when using bird-name labels? The answer is that in the present experiment the associations between blob-form and bird-label are arbitrary, whereas in Arguin et al. blob-form corresponded to the shape of these real-world objects. From the Arguin et al. study we know that ELM retains some knowledge of the visual properties of fruits and vegetables. The problem is that this knowledge is not exhaustive. Thus, ELM knows that a carrot is long and thin; he just isn't sure whether it is tapered or not, or whether it is curved. As such, ELM will confuse carrots, cucumbers, and bananas, but never confuses carrots with oranges or pears. Thus, when identifying the ET unbent set using fruit and vegetable labels, ELM correctly segments this set into thick shapes (melon and pear) and thin shapes (carrot, and cucumber). It then becomes possible for him to focus, in a subsequent step, on just a single dimension (in this case, just tapering) to identify each of the four exemplars correctly. After a few reminder trials, this subsequent step can be done flawlessly, but it takes time. The key is in the Arguin et al. study, ELM *already knows* that a carrot is long and thin (he just has to focus on the second diagnostic dimension). For the arbitrary blob–bird label pairings for each shape he must learn to extract values on both elongation and tapering; a situation that leads to more numerous errors. Unlike these semantically close sets (birds in the present study, fruits and vegetables in Arguin et al.), when semantically disparate concepts are applied to the blobs, ELM extracts values across multiple shape dimensions and is able to identify these objects in a quick and error-free manner.

Testing the Biological vs. Nonbiological Distinction

Although the preceding data are attributed to the interaction between shape set dimensionality and semantic proximity, it must be noted in both Arguin et al., and in Experiment 2 above, all of the semantically disparate sets consist of nonbiological objects and all the semantically close sets consist of biological objects. As previously noted, some authors have postulated that biological and nonbiological objects are each processed by specialised subsystems (e.g. Silveri et al., 1992). We are now in the position to test this assertion directly, without confounding semantic proximity and object structure.

EXPERIMENT 3

In Experiment 3, we pit the strong version of the biological-nonbiological hypothesis against the abnormal exemplar-node receptive field hypothesis, by assessing shape identification ability when shapes are paired with *man-made objects* that are semantically close to one another.

For real-world objects ELM shows the classic pattern of category-specific visual agnosia in which biological objects are impaired and most nonbiological objects are identified. According to the separate subsystems hypothesis it is conjectured that there is damage to the biological object recognition system, but not to the nonbiological system. If nonbiological objects are processed using a system that is somehow separate from the damaged system which processes biological objects, then ELM should not show recognition deficits for a set of nonbiological objects irrespective of their semantic proximity. If, on the other hand, both biological and nonbiological objects are processed by the same object recognition system, and problems arise when objects both share values on critical shape dimensions and are semantically close, then similar recognition difficulties would be predicted for biological and man-made objects.

Method

Subject

Testing was conducted only on ELM. The normal participants' often error-free performance in Experiments 1 and 2 precluded statistical analysis and served only to illustrate the ease with which these tasks can be completed by a healthy participant. (The manner in which object confusions in the memories of normal participants are affected by semantics can be found in Experiment 6, following a more complete assessment of ELM's object identification capabilities.)

Materials

Stimuli. The shapes in the conjunction set ET unbent were randomly paired with digitised recordings of a banjo, guitar, bass, or violin. (The ET unbent set was chosen to allow comparisons with his Experiment 2 performance, where this set was paired with either bird names or unrelated artefact labels.)

In addition shapes in the single dimension set E were randomly paired with digitised recordings of a banjo, guitar, bass, or violin.

Procedure. The same procedure was used as in Experiment 2.

Results

Conjunction Sets

Error rates for the ET unbent conjunction set were 60.41% for block 1 and 56.25% for block 2. In Experiment 2, which used the same shapes but paired these blobs to the sounds of a robin, crow, owl, and turkey, ELM had error rates of 50% on block 1 and 50% on block 2. Again in Experiment 2, the same shapes were paired with sounds of a saw, tennis-racquet, photocopier, and helicopter ELM's errors were 24.0% on block 1 and 0.0% on block 2.

Performance was not significantly different between the four bird song and four stringed instrument conditions ($\chi^2 = 1.33$, n.s.) for block 1 and ($\chi^2 = 0.62$, n.s.) for block 2. Performance was significantly poorer for the four stringed instruments relative to the four unrelated objects for both block 1 ($\chi^2 = 15.00$, $P < .001$) and block 2 ($\chi^2 = 56.00$, $P < .001$) respectively.

Single Dimension Set

Error rates for the E set were 30.00% for block 1 and 18.8% for block 2. In Experiment 2, which used the single dimension set B paired with robin, crow, owl, and turkey, ELM had error rates of 13.54% for block 1 and 10.42% on block 2. Also in Experiment 2, the E set was paired with sounds of a saw, tennis-racquet, photocopier, and helicopter. For this set ELM's errors were 16.00% for block 1 and 1.04% for block 2.

Performance was not significantly different between the four bird song and four stringed instrument conditions ($\chi^2 = 6.22$, n.s.) for block 1 and ($\chi^2 = 2.403$, n.s.) for block 2. Performance was also not significantly different between the four stringed instruments and the unrelated artefact sound on block 1 ($\chi^2 = 4.26$, n.s.) but was significantly different on block 2 (18.8% vs. 1.04%; $\chi^2 = 15.89$, $P < .01$). Conjunction and single dimension set performance for unrelated artefacts, birds, and stringed instruments are presented in Fig. 4.

Discussion

When shapes sharing values on critical shape dimensions are paired with experiences designed to elicit semantically similar object labels, recognition deficits occur in ELM. When exactly the same shapes are paired with experiences designed to elicit semantically disparate labels, however, identification performance improved dramatically. The finding that both biological and nonbiological object labels yield this pattern of results contradicts the notion of two separate object recognition systems subserving biological and nonbiological objects. Rather, this pattern favours a single object recognition system in which the interaction between multidimensional shape sets and semantic similarity conspires to create recognition problems for category-specific agnosic patients like ELM.

Once again, care was taken to ensure that ELM could distinguish flawlessly among the sound recordings from same and different categories prior to testing. It could still be argued, however, that the semantically close sounds were all produced using essentially the same mechanisms (vibrating strings, or the throats of birds or animals), whereas for semantically disparate experiences, sounds were produced in different ways (e.g. the scanner in the photocopier, and the rotating blade of the helicopter). Thus, it could be argued that the pattern of obtained results has nothing to do with the semantic proximity of the labels elicited by the sounds but is entirely due to the greater acoustic similarity of semantically similar sounds. That is, ELM had a harder time mapping the sound of a guitar or bass than sounds of a photocopier and helicopter to blobs, not

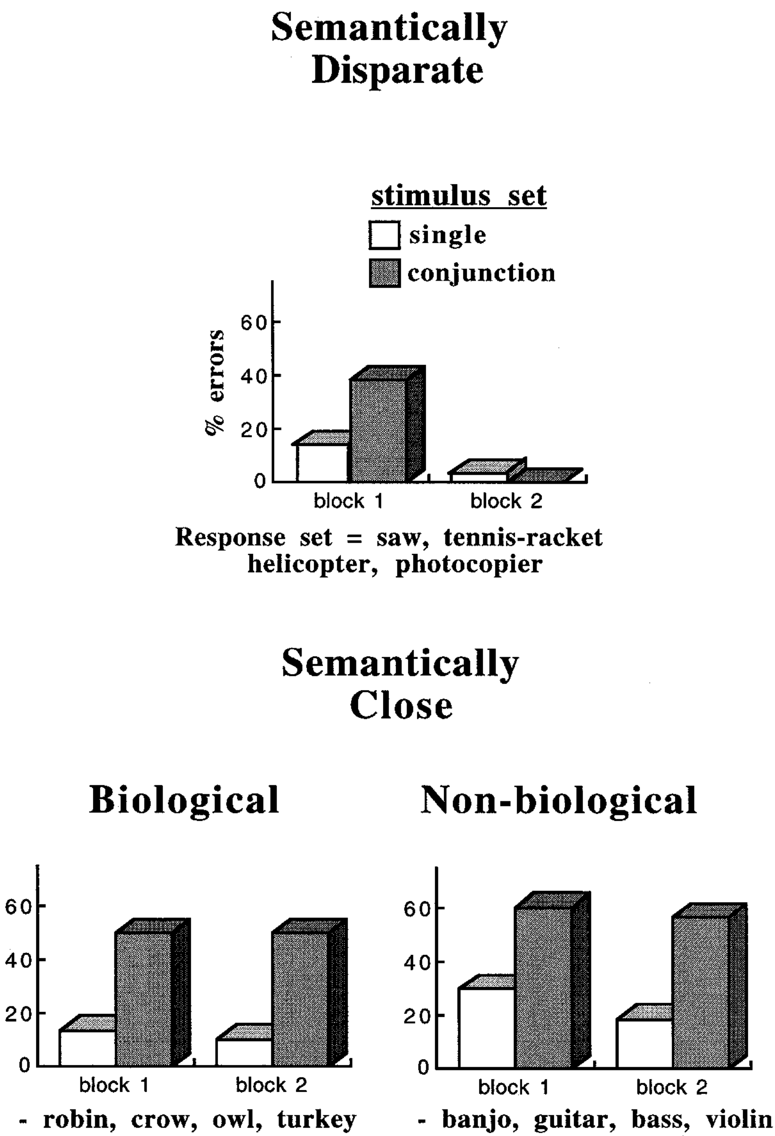

FIG. 4. ELM's block 1 and 2 error percentages for single dimension and conjunction sets when shapes were paired with semantically disparate concepts or with semantically close concepts belonging to either biological or nonbiological categories (Experiment 3).

because the stringed instruments are semantically closer, but merely because their sounds were more similar.

Such an argument might be considered untenable given that ELM identified the sounds flawlessly in pretesting. A different potential source of concern, however, is that for certain categories of sounds there is more than one way to differentiate between members of a given set. Pilot testing revealed, for example, that the digitised recordings of a car, bus, truck, and train (four vehicles) were primarily distinguished by hearing the gears changing in the car, the opening and closing of the bus door, the hydraulic brakes of the truck, and the whistle of the train. Thus, for certain sets, semantically close labels (car, bus, truck, and train) could have been derived from semantically disparate labels: gears, door, brakes, and whistle.

The simplest means of circumventing both of these potential artefacts is by dispensing with sounds and simply pairing shapes directly with verbal labels. Although originally it was thought that having ELM generate labels based on sound experiences might enable him to encode and remember the verbal labels better, according to Estes (1994), simply storing the categorical label itself in the memory array ensures that all of the accompanying categorical attributes will also be stored, and contribute to the overall similarity between exemplars.

EXPERIMENT 4

In Experiment 4, exemplars from single dimension and conjunction sets will be paired with verbal labels instead of digitised sound recordings. This allows us to retest the biological vs. nonbiological distinction against the exemplar-node receptive field deficit hypothesis in an even more rigorous fashion. In Experiment 4, names of man-made artefacts that are close in semantic proximity, and names of biological concepts that are semantically disparate, were mapped to single dimension and conjunction shape sets. If the biological vs. nonbiological distinction is the crucial factor in CSVA, then conjunction shapes paired with biological objects will not be identified, and conjunction shapes paired with artefacts will be identified, irrespective of the semantic distance between the concepts associated with the shapes. If, on the other hand, the receptive field deficit hypothesis is correct, then the biological vs. nonbiological dimension should be irrelevant. As long as shape sets share multiple diagnostic shape dimensions *and* bear labels that are semantically related, identification problems are predicted to ensue.

Materials and Methods

Stimuli. The following sets were used:

1. The semantically disparate words “shark”, “rose”, “apple”, “hummingbird” were paired with conjunction set (BE tapered), and single dimension set B.
2. The semantically close words “hammer”, “saw”, “wrench”, and “screwdriver” were paired with the conjunction set (BE not tapered) and single dimension set B.
3. The semantically disparate words “plate”, “door”, “stapler”, “kite” were paired with the conjunction set (ET bent) and single dimension set T.
4. The semantically close words “mustang”, “corvette”, “jeep”, and “cadillac” were paired with the conjunction set (ET bent) and single dimension set T, which are the same shapes as SET 3.

Procedure. The same procedure was used as in Experiment 3. The only exception was that digitised recordings of object sounds (e.g. sound of a saw cutting wood) were replaced by digitised recordings of the spoken verbal labels (e.g. the word “saw”).

Results

1. *Conjunction sets*

a. *(Semantically disparate):* For conjunction sets with semantically disparate labels, error rates ranged from 18% to 41% on block 1, and from 0% to 3% on block 2. These results are depicted by the dark bars in sets 1 and 3 of Fig. 5. Combining the semantically disparate sets, error rates were 29.2% on block 1 and 1.5% on block 2.

b. *(Semantically close):* Error rates ranged from 70% to 71% on block 1 and 66% to 21% on block 2. These sets are depicted by the dark bars in sets 2 and 4 of Fig. 5. Combining results from the two semantically close sets error rates were 70.3% on block 1 and 43.23% on block 2.

Using the combined data sets and comparing semantically close and disparate sets, the semantically close sets led to significantly higher error rates on both block 1 (70.3% vs. 29.17%, $\chi^2 = 32.67$, $P < .001$), and block 2 (43.23% and 1.50%, $\chi^2 = 74.41$, $P < .0001$).

2. *Single dimension sets*

a. *(Semantically disparate):* For single dimension sets with semantically disparate labels, error rates ranged from 8.33% to 21.88% on block 1, and from 10.41% to 15.62% on block 2. These error rates are depicted by the light bars in sets 1 and 3 of Fig. 5. Combining these semantically disparate single dimension sets, errors were 29/192 (15.10%) on block 1 and 25/192 (13.02%) on block 2.

b. *(Semantically close):* Error rates ranged from 21.89% to 20.83% on the first block of 96 trials and between 10.41% and 16.67% on the second block of 96 test trials. These error rates are depicted by the white bars in sets 2 and 4 of

Fig. 5. Combining these semantically similar single dimension sets, error rates were 21.35% on block 1 and 13.54% on block 2.

Using the combined data sets and comparing semantically close set, to semantically disparate sets, no significant differences were obtained for single dimension sets. Error rates were not significantly different between semantically disparate and close sets on block 1 (15.10% vs. 21.35%; $\chi^2 = 1.07$, n.s.) or between disparate and close sets on block 2 performance (13.02% and 13.54%, $\chi^2 = 0.01$, n.s.)

3. *Identical shapes comparison.* A comparison of sets 3 and 4, in which the same conjunction shapes were used for both semantically similar and disparate concepts, revealed equivalent performance after the first block (40.63% disparate errors vs. 31.2% close errors, $\chi^2 = 1.24$, n.s.) but significantly fewer errors for the disparate set for the second block (3.1% vs. 25% close errors, $\chi^2 = 16.33$, $P < .01$).

ELM's single dimension and conjunction set performance for identical shapes associated with verbal labels of semantically close and disparate proximity are presented in sets 3 and 4 of Fig. 5.

Discussion

Experiment 4 replicates and extends the findings of Experiments 2 and 3. It rules out the possibility that the pattern of results produced by these previous experiments was attributable to an artefact involving sound–shape pairings in which semantically similar sounds were more acoustically similar than semantically unrelated sounds. Using verbal labels, where acoustic similarity is not an issue, ELM once again displayed profound recognition difficulties only when conjunction sets were paired with semantically similar concepts.

Once again, the biological-nonbiological distinction does not appear to be the relevant variable in these experiments, as ELM has no problem recognising a conjunction set of blobs mapped to four biological concepts (shark, rose, apple, and hummingbird) that are semantically disparate, but has a great deal of difficulty identifying a conjunction set of blobs mapped to nonbiological concepts (corvette, mustang, cadillac, and jeep) that are semantically similar.

As in Experiments 2 and 3, the powerful effect of semantic proximity is underscored by this pattern of results occurring even when the same shapes serve for semantically close and semantically disparate sets. By using the exact same shapes, the visual structure of objects within a set is held constant while semantic proximity is independently modified—a situation that provides the most rigorous test of the shared feature by semantic proximity interaction hypothesis. Like Experiments 2 and 3, by the end of the second block performance was nearly flawless for the semantically disparate conjunction sets, but was still quite poor when these same shapes were paired with semantically close concepts.

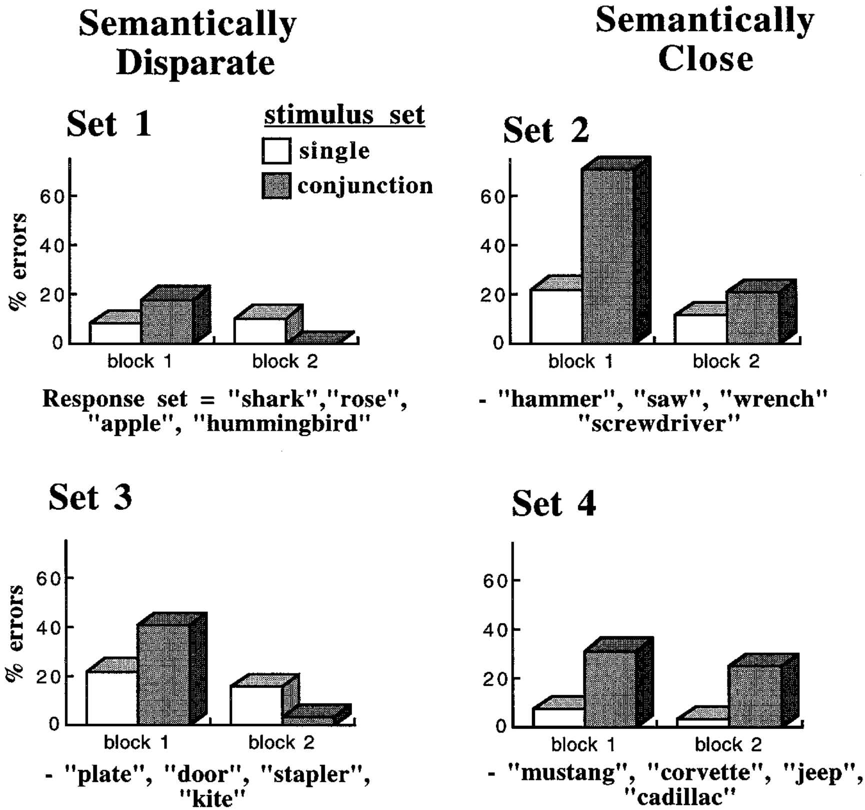

FIG. 5. ELM's block 1 and 2 error percentages for single dimension and conjunction sets when shapes were paired with semantically close or disparate concepts evoked through verbal labels (Experiment 4).

Thus far, we have sampled sets of shapes associated with semantic labels that were dichotomously either semantically close or disparate. Exemplar models such as our variant of ALCOVE propose that objects are stored in multidimensional space at coordinates based on visual and semantic dimensions. In such a psychological space the proximity between two objects will depend on the number of overlapping visual and semantic attributes shared by these objects. Since the number of shared attributes can be considered a continuous, as opposed to a dichotomous, variable, it is of interest to compare categories that intuitively vary in the number of overlapping attributes. Thus, although car, bus, truck, and train present some degree of attribute overlap, one can guess that these categories are not as semantically close as four types of car (e.g. mustang, corvette, jeep, and cadillac), which in turn are less semantically proximate than four sports cars (e.g. mustang, corvette, trans am, and camaro).

EXPERIMENT 5

Experiment 5 tested ELM's conjunction and single dimension shape identification performance across a range of intuitively derived semantic proximities. These new sets were combined with the previously acquired data from experiments 2 through 4. Semantic proximity ratings were then gathered from normals, and these empirically derived proximities were used to formally measure the association between ELM's semantic proximity and conjunction set performance.

Method

Subjects

ELM participated in the shape identification paradigm. For the semantic rankings normal participants were recruited. These participants were 31 university students enrolled in a third-year cognition course at the University of Victoria. Participants participated in the study in fulfilment of a course requirement.

Materials

Shape identification stimuli. The following word quadruplets were paired with the following conjunction and single dimension sets:

1. The words "hummingbird", "lion", "wasp", and "frog" were paired with the conjunction set (BT elongated), and single dimension set E.
2. The words "submarine", "metro", "airplane", and "bus" were also paired with the conjunction set (BT elongated) and single dimension set E.
3. The words "corvette", "trans am", "mustang", and "camaro" were paired with the conjunction set (BT not elongated) and single dimension set T.
4. The words "robin", "crow", "cardinal", and "bluejay" were also paired with the conjunction set (BT not elongated) and single dimension set T.
5. The words "cup", "bowl", "glass", "vase" were paired with the conjunction set (BE tapered) and single dimension set B.
6. The words "car", "bus", "truck", and "train", were paired with the conjunction set (BE tapered) and single dimension set E.

Procedure. The same procedure as in Experiment 4 was used to test ELM.

Semantic rankings from normals. Participants were given the 15 sets of concepts depicted in Table 1. The 15 quadruplets were presented in random order and participants were instructed to read all 15 quadruplets and then to rank order them in terms of how similar the members within the quadruplets were. Rankings were to be ordered from least similar = 1 to most similar = 15.

TABLE 1

Total, Block 1, and Block 2 Error Percentages, for Single Dimension (1D) and Conjunction Shape (2D) Sets as a Function of Semantic Proximity Rankings

			Errors (%)				Errors (%)		
Concepts	*Rank*	*2D Sets*	*Block 1*	*Block 2*	*Total*	*1D Sets*	*Block 1*	*Block 2*	*Total*
Leaf blower, glass motorcycle, telephone	2.10	BE 0 tap	38.54	0.00	19.27	Ben	37.50	6.00	21.88
Saw, tennis-raquet, helicopter, photocopier	2.40	ET ben	30.21	0.00	15.10	Elong	16.00	1.04	8.33
(Replication of above in Exp. 2)[a]		ET 0 ben	24.00	0.00	12.00				
Kite, door, plate, stapler	3.00	ET ben	40.63	3.13	21.88	Tap	21.88	15.63	18.75
Shark, rose, apple, hummingbird	3.14	BE 0 tap	17.71	0.00	8.85	Ben	8.33	10.42	9.38
Lion, wasp, frog, hummingbird	5.61	BT elong	5.21	0.00	2.60	Tap	1.04	1.04	1.04
Dog, horse, wolf, elephant	7.61	BE tap	67.71	32.29	50.00	Elong	37.50	18.75	28.13
Submarine, metro, airplane, bus	8.50	BT elong	36.46	20.83	28.65	Elong	11.46	0.00	5.73
Robin, crow, owl, turkey	9.20	ET 0 ben	50.00	50.00	50.00	Ben	13.54	10.42	11.98
Cup, bowl, glass, vase	9.25	BE tap	71.00	66.00	68.23	Ben	10.00	3.10	7.29
Car, bus, truck, train	9.71	BE tap	55.21	22.92	39.06	Elong	8.33	2.08	5.21
Jeep, cadillac, corvette, mustang	9.96	ET ben	31.25	25.00	28.13	Tap	7.29	3.13	5.21
Hammer, saw, wrench, screwdriver	11.07	BE tap	70.83	20.83	45.83	Ben	21.88	11.46	16.67
Crow, robin, cardinal, blue jay	11.89	BT 0 elong	70.00	52.08	60.94	Tap	30.21	7.29	18.75
Banjo, guitar, bass, violin	12.40	ET 0 ben	60.42	56.25	58.33	Elong	30.00	18.80	24.48
Trans am, mustang, camaro, corvette	14.21	BT 0 elong	69.79	65.63	67.71	Tap	20.83	16.67	18.75

[a]This replication was not entered in the correlation.

The following quadruplets were given as examples: "If the word quadruplet was (daisy, rose, chrysanthemum, and tulip) it should get a "high" (near 15) ranking since these are all flowers, all grow in Canadian gardens, all have long stems, all have green stems, petals, etc."

"On the other hand if the quadruplet was (building, fig leaf, computer, needle) this quadruplet would get a low number (close to 1) because these items more or less have nothing to do with one another."

Participants were then instructed to ". . . use everything you know about the items when deciding on the rankings (what the items look like, what the items do, what the items are used for, etc.)"

Finally, participants were advised to ensure that they only had one ranking per quadruplet and that they had as rankings the numbers 1 to 15 inclusive when they were finished.

Results

Despite explicit instructions, three normal participants reversed the order of the rankings (gave low rankings to semantically similar items, and high rankings to disparate items). These data were removed from further analysis.

Data from the remaining 28 participants were used to compute semantic proximity values for each of the 15 quadruplets. These values were the means of the 28 rankings obtained for each quadruplet.

ELM's conjunction and single dimension shape set identification performance, along with semantic proximity values obtained from normals for each word quadruplet, are given in Table 1. For comparison with previous experiments, both total error rates and block 2 error rates are given in this table.

The correlation between semantic proximity values and ELM's shape identification performance was significant and strong for both total ($r = .81$, $P < .01$) and block 2 performance ($r = .84$, $P < .01$). Associations between semantic proximity values and single dimension set performance were negligible ($r = .21$, n.s., and $r = .06$, n.s. for total and block 2 error rates, respectively). The relationship between normal semantic rankings and ELM's block 2 performance is depicted in Fig. 6.

Finally, the quadruplets' mean word frequencies (Kucera & Francis, 1967) did not correlate with either total conjunction set error performance ($r = .016$, n.s.) or block 2 performance ($r = -.13$, n.s.)

Discussion

Experiment 5 replicates and extends the findings of Experiments 2 through 4 using a larger data set and more powerful parametric statistics. Before expanding on these findings we will dispense with some criticisms that could be brought to bear on this experiment.

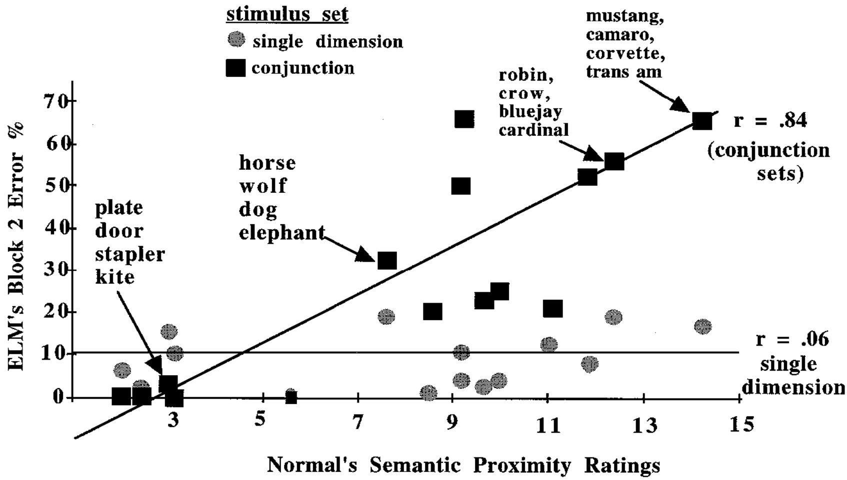

FIG. 6. ELM's block 2 error rates for single dimension and conjunction sets as a function of semantic proximity. Date are presented for ELM's block 2 errors (Experiment 5). The flat line of best fit reflects the minimal correlation between semantic proximity and single dimension sets. The sloped line of best fit, and close proximity of conjunction set points to this line, reflect the substantial correlation between semantic proximity and conjunction set performance.

First, we combined quadruplets from experiments that used sounds to portray concepts with quadruplets using verbal labels to portray concepts. Although undesirable, the alternative would be to retest ELM on all of these quadruplets using labels. Rather than trying ELM's (Herculean) patience, we chose to combine the two types of quadrants and accept the fact that different procedures could potentially contribute to unexplained variance, which could reduce the correlation between semantic proximity ratings and ELM's identification performance.

Second, although we controlled for the semantic proximity of the blob labels, we did not control for whether the real-world objects denoted by these labels were recognisable to ELM. Thus, it could be postulated that ELM only has trouble with conjunction sets of blobs with labels denoting objects that were unrecognisable for him. The conjunction set labeled with bowl, cup, glass, and vase dispenses with this notion. Despite his ability to recognise these real-world objects, his overall conjunction set performance on this set was the poorest (66% block 2 errors) of the 15 quadruplets tested. This contrasts with his nearly flawless performance (3.1% block 2 errors) on a conjunction set using kite, plate, stapler, door—four objects that he also can recognise. These vast performance differences suggest that whether or not he can recognise the objects denoted by the labels we attached to the blobs is of little importance to his blob-identification performance.

Finally we did not control for whether object labels consisted of base-level or subordinate-level terms. Indeed, certain subordinate-level terms were chosen especially for their semantic relatedness (e.g. mustang, camaro, trans am, and corvette). It is likely, however, that conjunction set performance was determined by the semantic proximity of these objects rather than by whether they were base-level or subordinate entities. ELM does equally poorly for semantically related base-level terms (bowl, cup, glass, vase = 66% block 2 errors) as he does for semantically related subordinate-level terms (mustang, camaro, trans am, corvette = 66% block 2 errors).

What is shown conclusively in Experiment 5 is that ELM's blob identification depends on both the dimensionality of the shape sets used and the semantic proximity of the concepts to which these blobs were associated. Experiment 5 once again demonstrated that for single dimension sets, ELM can disambiguate exemplars by gating irrelevant dimensions and increasing the attentional dimension strengths of single relevant dimensions. For these sets the effect of semantic proximity is minimal: four sports-car blobs are identified as easily as four blobs mapped to unrelated members of the animal kingdom (wasp, lion, hummingbird, frog).

It can be assumed that ELM's ability to distinguish items from shape sets varying along a single dimension transcends the confines of the laboratory. In real life ELM does not confuse cup, bowl, glass, and vase despite their intuitive semantic closeness. This is possibly because in real life ELM has optimised his ability to differentiate these exemplars by increasing the attentional dimension strength for elongation (a vase is taller than a glass, which is taller than a cup, which is taller than a bowl). Importantly, in the computer-generated shape paradigm, ELM can be made to confuse these semantically proximate entities if they are artificially forced to form a conjunction set rather than a single dimension set.

Unlike single dimension sets, when members within a set share values on critical shape dimensions, semantic proximity becomes a reliable and robust determinant of identification performance. By showing strong, significant correlations between semantic proximity values and conjunction set performance, Experiment 5 extends the findings of previous experiments by showing that it is not just the presence or absence of semantic proximity, but rather the degree of proximity, which is crucial to ELM's shape identification. Thus, although four vehicles (car, bus, truck, and train) pose some problems for ELM because they are somewhat proximate, very close semantic sets such as four birds, four stringed instruments, or four sports cars significantly exacerbate ELM's identification problems.

EXPERIMENT 6

In the experiments presented thus far two fundamental assumptions have been made. The first is that ELM has exemplar nodes with abnormally wide receptive fields. This causes recognition problems only for certain classes of objects because of the second fundamental assumption—namely that the manner in which we store objects in memory is based on psychological as opposed to just visual distance. According to this principle, the more visually and semantically similar two objects are, the greater is their psychological similarity, and (because of abnormally large receptive fields) the more likely it is that ELM will confuse them.

Experiments 1–5 support the contention the ELM's object identification errors are constrained by this psychological distance principle. It should, however, be theoretically possible to use the ELM paradigm to demonstrate that healthy adults (with presumably normal-sized receptive fields) also store objects according to this psychological distance principle. To achieve this, one must show that normals, like ELM, make more object confusions when blobs are labeled using semantically close concepts than they make when the same blobs are associated with semantically disparate concepts. The nearly flawless performance of the healthy participant in Experiments 2–3 indicates that mapping four sounds or labels to four shapes runs the risk of generating ceiling effects. Thus, in testing normals, set size was increased to six blobs and six labels in the hope of generating enough object confusions to enable the ramifications of the psychological distance principle to emerge. In order to maximise the effects of semantic proximity, in the semantically close condition participants were asked to map bird names to blobs. Birds are one of the visually similar and semantically similar sets of objects known to man. In the unrelated condition exemplars were chosen from six completely different semantic categories (insect, amphibian, mammal, tool, instrument, vehicle).

It was predicted that blob-labeling errors would increase significantly when related labels replaced semantically unrelated labels in the blob-labeling task. However, this effect of semantic proximity was predicted to be smaller than that noted for ELM, who should show more pronounced differences between the semantically related and unrelated conditions because of abnormally wide receptive fields. (Unfortunately, generating six exemplars varying along a single dimension caused exemplars to be too close perceptually to be useable in this paradigm. A plethora of research suggests, however, that single dimension sets are classified better than sets of objects defined by multiple visual dimensions: see Kruschke, 1992, for a review.)

Method

Subjects

Eighteen elderly participants were tested. Of these, four were unable to perform above chance on either condition and were excluded from further analyses. The 14 remaining participants were somewhat older than ELM, ranging in age from 66 to 95 years old (mean age = 79.84 years). Participants were all living independently, mentally healthy, with no subjective memory complaints.

Materials

Shape stimuli. Six of the eight blobs used to form conjunction sets were employed. These were the rightmost blobs from the elongated and unelongated BT sets of Fig. 1 (i.e. the BT sets excluding the cigar-shaped and watermelon-shaped blobs). In terms of visual similarity each of the six blobs within this set shared two features with at least one other exemplar.

Labels. The following sets of labels, matched for word frequency, were applied to the six shapes:

1. The words "robin", "sparrow", "crow", "cardinal", "bluejay", "swallow" were paired with the conjunction set described earlier.
2. The words "frog", "tiger", "wasp", "carriage", "wrench", "banjo" were paired to these same blobs.

Procedure

Learning trials. On learning trials one of the six shapes were presented for 1000msec along with its acoustically presented name. Following a 1000msec interstimulus interval, a second shape was presented for 1000msec accompanied by its name. Six such learning trials were presented with each shape–name combination presented once.

Test trials. Following six learning trials, six test trials were presented where shapes were unaccompanied by their names. Participants attempted to give the names associated with the shapes. The pattern of 6 learning and 6 test trials was repeated 12 times for a total of 72 learning and test trials. The relevant data are the number of errors made over 72 test trials.

Participants were tested on separate sessions conducted at least a week apart. Session order was counterbalanced across participants.

Results

Semantically close error rates were significantly higher (mean = 53.5%) than the error rates for the unrelated labels (mean = 36.8%) [dependent $t(12) = 3.40$, $P < .01$].

For the blob quadruplets mapped to bird names ELM's total error percentage was 60.94%. For quadruplets mapped to unrelated (both biological and nonbiological) labels the best estimate of ELM's error rate is 13.75%. This value came from the percentage of errors ELM made on the five quadruplets with the lowest semantic proximity rankings (the five quadruplets at the top of Table 1)[4].

A two by two chi-square analysis of error percentages by ELM and normals for the semantically close and disparate conditions indicates significantly greater effects of semantic proximity (60.94% close vs. 13.75% disparate) for ELM than for healthy participants (53.5% close vs. 36.8% disparate errors) ($\chi^2 = 9.86$, $P < .01$).

For the semantically close condition, healthy participants (53.5% errors) performed similarly to ELM (60.94% errors), ($\chi^2 = 0.481$, n.s.) For the semantically disparate conditions, ELM's performance (13.75% errors) was significantly better than that of normals (36.8% errors) ($\chi^2 = 10.82$, $P < .001$).

Discussion

The performance of healthy participants in a more difficult version of the ELM paradigm provides evidence that objects are retrieved from memory according to a psychological distance principle. Objects that share only visual attributes are further apart in psychological space and are less confusable than objects that are stored close together by virtue of sharing both visual *and* semantic attributes. Nevertheless, normals do not show nearly as marked a discrepancy between the semantically close and disparate label conditions as ELM did.

In this experiment normal performance was moved away from ceiling by assessing their ability to attach six blobs to six labels. For shapes mapped to semantically disparate labels, this increase in set size caused normals to make significantly more errors than ELM. For exactly the same shapes mapped to semantically close labels, however, ELM's performance with four blobs rose to levels that were ordinally (albeit not statistically) higher than normals' error rates with six blobs. ELM's poorer performance, despite using a smaller set size, is consistent with the notion that he has preferential difficulties disambiguating objects that share visual *and* semantic properties.

GENERAL DISCUSSION

The utility of the shape labeling task in understanding object recognition depends on its relationship to veridical object recognition. We propose that in their most important aspects, the two processes are analogous. In the versions of the ELM paradigm used in Experiments 4–6, on learning trials a blob was

[4]This is a conservative estimate. The condition that ELM completed which was most comparable to the control subjects' unrelated condition was that using labels of an animal, an amphibian, a bird, and an insect mapped to the BT elongated shapes. In this condition ELM made only 2.5% total errors.

accompanied by an auditorially presented name which automatically makes contact with semantic representations corresponding to that name. In order to succeed on test trials subjects must somehow learn to map the shape to the name that accompanied it on learning trials. To do this they may either go directly from the form of the blob to phonology (the blob's "name"), or along a route that travels through semantics. ELM's strong effect of semantic proximity suggests that the pathway taken runs through semantics.

Thus, the blobs in the ELM paradigm can be thought of as descriptions in shape space, just like the form of real objects. Like real objects, these points in shape space make contact with relevant points in semantic space, which in turn make contact with phonology. These mappings probably all take place in cascade fashion (Humphreys et al., 1988). The central difference between veridical object recognition and recognition in the ELM paradigm, therefore, is that in veridical object recognition these mappings between form, semantics, and phonology already exist, whereas in the ELM paradigm these mappings must be instantiated.

Where this paradigm excels in relation to veridical object recognition or confrontation naming of standardised line drawings is in its ability to decouple object form from object meaning. By having exactly the same shapes standing for both semantically close and disparate objects one can assess directly the ramifications of semantic proximity on object recognition unconfounded by visual proximity. By applying the same labels to visually similar or visually distinct blob sets, for the first time in Neuropsychology one can directly assess the ramifications of visual proximity unconfounded by semantic proximity. With ELM we showed that for sets of objects sharing values on multiple shape dimensions, performance depended completely on semantic proximity. For semantically close labels, performance was poor; for semantically disparate items mapped to exactly the same shapes, performance was near ceiling. In addition, we repeatedly showed that the biological vs. nonbiological distinction is of minimal importance relative to the pairing of visually similar shapes with semantically similar concepts.

Thus, the present series of experiments provides the first full empirical evidence that the identification problems in at least some forms of category-specific visual agnosia can result from an interaction between the shape set dimensionality and semantic proximity of the objects being identified. This interaction can be interpreted within the context of an exemplar model of categorisation and identification that makes three fundamental assumptions. The first is that ELM has exemplar nodes with abnormally wide receptive fields. The second assumption is that psychologically similar objects are stored close together in multidimensional psychological space. The third assumption is that psychological distance is derived from both visual features and semantic features. Together these assumptions demonstrate why ELM shows CSVA—the abnormally wide receptive fields of contiguously stored objects will overlap making visually similar *and* semantically similar objects extremely difficult to recognise. An important exception to this general principle involves sets of objects in which the exemplars differ along a single shape dimension. In this case ELM can compensate for abnormally wide receptive fields by optimising the dimensional attention strength associated with these single relevant dimensions.

This psychological distance view of category-specific agnosia contrasts with other models. Of these, one of the most influential is that proposed by Warrington and colleagues (Warrington & McCarthy, 1987, 1994; Warrington & Shallice, 1984). They postulate that semantics is parsed into knowledge concerning sensory properties (e.g. what the object looks like) and knowledge concerning function (what the object does). If the knowledge of sensory properties becomes damaged, biological objects become unrecognisable because discrimination among exemplars relies primarily on visual features. Artefacts can still be recognised because patients retain knowledge of their often unique functions. Thus, category-specificity has nothing to do with biological category membership per se, but rather, whether or not an object can reliably be identified according to its functions.

Recent evidence from PET studies of healthy individuals at least partially supports this view. Martin, Wiggs, Ungerleider, and Haxby (1996) used subtraction methodology to show that identifying animals draws upon ventral areas of the temporal lobes as well as the primary visual cortex. Such preferential visual cortical activation for animals but not tools was interpreted as a re-consultation of the animal's fine-grained visual features in order to arrive at base-level identification. When subjects identify tools, temporal lobe activation is accompanied by activation in the frontal lobes. Martin et al. attributed such frontal activity to activation of the cortical areas responsible for encoding knowledge about object function. Thus, as Warrington et al. suggest, the disambiguation of living things relies primarily on visual features, whereas the disambiguation of objects like tools involves their function.

Because ELM has identification problems primarily with living things, Warrington et al. would assume damage to his sensory knowledge subsystem, but intact knowledge concerning the functions of objects. This would allow ELM to recognise things like tools in everyday life. In the blob-labeling paradigm ELM had fewer problems learning to pair single dimension shapes to tool names (16.67% errors) than conjunction shapes (45.83% errors). Warrington et al. might argue that forcing ELM to disambiguate exemplars sharing multiple visual features draws upon his damaged sensory knowledge system and leads to numerous errors. Warrington et al. might also note that conjunction set errors for shapes with tool labels (20.83% block 2 errors) were substantially less than conjunction set errors for shapes with bird labels (52.08% block 2 errors)—a finding consistent with the notion that the unique functions associated with tools aids in their identification. Where Warrington and colleagues' theory encounters problems is explaining why ELM made *no errors* in block 2 on conjunction sets labeled using “shark”, “rose”, “apple”, and “hummingbird” or “lion”, “wasp”, “frog”, and “hummingbird”—biological objects that, like birds, have *no* unique functions for man. Taken together, one must conclude that for ELM, the role of functional knowledge of objects may serve to increase semantic distance somewhat but it cannot account entirely for his different patterns of performance for single and conjunction sets of objects. The errorless performance of ELM for conjunction sets labeled using functionless biological objects indicates that semantic distance plays a more crucial role than object function in determining object identification performance.

Another difficulty in assuming that artefacts can be recognised by their functions involves the problematic category of musical instruments. A guitar has a salient function for man. If objects can be recognised by their functions, why, then, should objects like a guitar pose identification problems for patients who otherwise have difficulties predominantly with biological objects. The answer once again involves psychological distance. If nonbiological objects like tools (e.g. saw and hammer) have different functions, this would serve to increase the semantic distance between the exemplars comprising these categories, thereby making these objects easier to recognise. For objects like a guitar, however, their function is quite similar to other exemplars within the subcategory of stringed musical instruments. Violin, guitar, and banjo all have similar forms, similar functions, and (at least concerning the left hand) require similar kinesthetic movements. This overlapping of structural, functional and kinesthetic attributes may lessen semantic distance and lead to recognition problems for these items.

It should be noted, however, that our exemplar-node receptive field hypothesis cannot account for all forms of category-specific agnosia. It cannot account for the less common form of CSVA in which patients show recognition deficits for man-made artefacts but a sparing of biological objects (Hillis & Caramazza, 1991; Sacchett & Humphreys, 1992; Warrington & McCarthy, 1983, 1987).

For patients who show the more prevalent pattern of category-specific visual agnosia (recognition problems with biological categories and a sparing of artefacts), a deficit that leads to a strong interaction between semantics and shape set dimensionality must send a cautionary note to researchers who would look exclusively to the level of semantics for an explanation of this form of CSVA. These findings would suggest that extreme care be taken in ruling out shape processing deficits before attributing the cause of CSVA solely to a semantic deficit in a given patient. As has repeatedly been shown in Experiments 2–5, shape processing deficits can interact with semantics in spectacular fashion.

The shape identification performance of ELM also has implications for general theories of object recognition. It provides clear evidence that semantics can modulate shape identification. Thus, for exemplar models like ALCOVE, semantic must be involved in determining the position of the hidden exemplar nodes in multidimensional psychological space. Although the purpose of this study was not to provide a fully specified model capable of object recognition, it does suggest that at least two key elements must be present in such a model if it is to account for ELM's identification performance. First, some form of dimensional gating mechanism like that in ALCOVE is necessary to account for ELM's performance differences between single dimension and conjunction shape sets. Second, a mechanism that allows both semantics and structural factors to enter into the calculation of exemplar similarity must be present in order to account for semantics modulating ELM's conjunction set identification performance.

ELM's propensity to confuse objects that are both visually and semantically similar is somewhat reminiscent of the "mixed" visual and semantic reading errors (e.g. reading "rat" as "cat") made by patients with deep dyslexia (Hinton & Shallice, 1991). In simulations of many (but not all) connectionist configurations, mixed errors are a prevalent form of deep dyslexic reading errors, and can arise from damage to a number of different model components (Hinton & Shallice, 1991; Plaut & Shallice, 1993).

For object identification, there may be an even greater propensity to confuse items that are visually and semantically related. This is because in reading, word form is only arbitrarily related to semantics (the word "CAT" does not look like the four-legged feline), and large attractor basins are proposed to overcome the problems distributed architectures have in mapping visually similar inputs (CAT, MAT) to disparate patterns of activations in semantic space. For visually presented objects, on the other hand, visually similar forms are often also semantically similar (e.g. the shapes of a robin and a crow are similar, as are their meanings). Distributed architectures require less learning and smaller connection strengths to map visually similar inputs to highly similar, but still discernibly different, patterns of activation in semantic space (Plaut & Shallice, 1992). As in deep dyslexia, however, a consequence of storing object representations in this distributed fashion would be that damage to this architecture would be likely to result in the propensity to confuse objects that are both visually and semantically related. Thus, whether accounting for the word reading errors of deep dyslexic patients, or object recognition problems of temporal lobe patients like ELM, if one assumes that knowledge is stored in a distributed architecture, then confusions among entities that are both visually and semantically similar are to be expected when this architecture becomes damaged.

Irrespective of cognitive neuropsychologists' view of whether object identification is best explained using exemplar models or more distributed architectures, the paradigm employed in these experiments offers an unprecedented ability to look directly at the influence of semantic relationships in category-specific visual agnosia. By allowing the same shapes to stand for concepts that differed in semantic proximity we have effectively held the structure of the objects constant, while independently manipulating their semantic properties. Using this paradigm we have repeatedly shown that shape identification in the category-specific visual agnosic patient ELM depends on the interaction of visual feature overlap and semantic proximity.

It can be postulated that in real life, most man-made objects are either visually dissimilar and hence pose no problems, or category members can be differentiated from one another using a single shape dimension (like bowl, cup, glass, and vase). Finally, nonbiological objects may pose fewer problems than biological objects because they have specific and often unique functions; a situation that might serve to increase the semantic distance between members of nonbiological categories. Exceptions are categories like makes of car and musical instruments, which have similar functions and pose problems for ELM and other category-specific agnosics because they probably require more than one crucial shape dimension for their disambiguation.

Unlike man-made artefacts, biological objects not only share a large number of semantic features, but also share a large number of visual features (all animals have heads, necks, trunks, and legs). Thus objects like fruits, vegetables, animals, birds, and insects pose the deadly combination of semantic proximity and shared values along critical shape dimensions, which precludes object recognition in at least some forms of category-specific visual agnosia.

Manuscript received 22 November 1996
Revised manuscript received 13 November 1997
Manuscript accepted 22 December 1997

REFERENCES

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

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

Damasio, A.R. (1990). Category-related recognition defects as a clue to the neural substrates of knowledge. *Trends in Neuroscience*, *13*, 95–98.

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

Etcoff, N.L., Freeman, R., & Cave, K.R. (1991). Can we lose memories of faces? Content specificity and awareness in a prosopagnosic. *Journal of Cognitive Neuroscience*, *3*, 25–41.

Farah, M.J., Hammond, K.M., Mehta, Z., & Ratcliff, G. (1989). Category-specificity and modality-specificity in semantic memory. *Neuropsychologia*, *27*, 193–200.

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., McMullen, P.A., & Meyer, M.M. (1991). Can recognition of living things be selectively impaired? *Neuropsychologia*, *29*, 185–193.

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

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

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

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

Humphreys, G.W., & Riddoch, M.J. (1987). On telling your fruit from your vegetables: A consideration of category-specific deficits after brain damage. *Trends in Neuroscience*, *10*, 145–148.

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

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

Kucera, H.F., & Francis, W.N. (1967). *Computational analysis of present-day American English*. Providence, RI: Brown University Press.

Martin, A., Wiggs, C.L., Ungerleider, L.G., & Haxby, J.V. (1996). Neural correlates of category-specific knowledge. *Nature*, *379*, 649–652.

Mauri, A., Daum, I., Sartori, G., Riesch, G., & Birbaumer, N. (1994). Category-specific semantic impairment in Alzheimer's disease and temporal lobe dysfunction: A comparative study. *Journal of Clinical and Experimental Neuropsychology*, *16*, 689–701.

McClelland, J.L. (1979). On the time relations of mental processes: An examination of systems of processes in cascade. *Psychological Review*, *86*, 287–330.

Mehta, Z., Newcombe, F., & DeHaan, E. (1992). Selective loss of imagery in a case of visual agnosia. *Neuropsychologia*, *30*, 645–655.

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

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

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

Ratcliff, G., & Newcombe, F. (1982). Object recognition: Some deductions from the clinical evidence. In A.W. Ellis (Ed.), *Normality and pathology in cognitive function*. Toronto: Academic.

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*, 105–132.

Sartori, G., Job, R., & Coltheart, M. (1992). The organization of object knowledge: Evidence from neuropsychology. In D.E. Meyer & S. Kornblum (Eds.), *Attention and performance XIV* (pp. 451–465). Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

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

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.

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

Snodgrass, J.G., & Vanderwart, M.A. (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.

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

Warrington, E.K., & McCarthy, R.A. (1987). Categories of knowledge. *Brain*, *110*, 1273–1296.
Warrington, E.K., & McCarthy, R.A. (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 impairments. *Brain*, *107*, 829–854.