

## **Shape Integration for Visual Object Recognition and Its Implication in Category-Specific Visual Agnosia**

Martin Arguin

*Département de psychologie and Groupe de recherche en neuropsychologie expérimentale, Université de Montréal, Montreal, Quebec, Canada*

*Laboratoire Théophile-Alajouanine, Centre Hospitalier Côte-des-Neiges, Montreal, Quebec, Canada*

Daniel Bub

*Laboratoire Théophile-Alajouanine, Centre Hospitalier Côte-des-Neiges, Montreal, Quebec, Canada*

*Department of Psychology, University of Victoria, Victoria, British Columbia, Canada*

Gregory Dudek

*Center for Intelligent Machines, McGill University, Montreal, Quebec, Canada*

A series of experiments was conducted on a patient (ELM) with bilateral inferior temporal lobe damage and category-specific visual agnosia in order to specify the nature of his functional impairment. In Experiment 1, ELM performed a task of picture/word matching that used line drawings of fruits and vegetables as stimuli. The pattern of confusions exhibited by the patient suggested a failure in processing the full range of shape features necessary for the unique specification of the target relative to other structurally related items. This hypothesis of a shape integration impairment was tested and verified by subsequent visual recognition experiments

---

Requests for reprints should be sent to Martin Arguin, Département de psychologie, Université de Montréal, C.P. 6128, Succ. Centre-ville, Montréal, Québec, H3C 3J7, Canada. E-mail: arguinm@ere.umontreal.ca

We are deeply indebted to ELM for his patient and lasting collaboration. Thanks also to Leen Werbrouck for her superb technical assistance. This paper has benefited from comments by Drs G.W. Humphreys, R.A. McCarthy, M.J. Riddoch, and by an anonymous reviewer on a previous version. This project was supported by a post-doctoral fellowship and an operating grant from the Medical Research Council of Canada to Martin Arguin, an MRC operating grant to Daniel Bub, and a grant from the Natural Sciences and Engineering Research Council of Canada to Gregory Dudek.

(Experiments 2–4), which used synthetic stimuli with shapes precisely defined on the dimensions of elongation, curvature, and tapering. Furthermore, it was determined (Experiment 5) that the integration deficit is specific to the *retrieval* of shape knowledge from memory and does not affect the *encoding* of the properties of visual stimuli. It is argued that these findings have critical implications for cognitive theories of visual object recognition and for an interpretation of the visual function of the inferior temporal cortex. Finally, it was shown that the patient's deficit for structural knowledge integration is modulated by the semantic properties of the objects (Experiment 6), thereby demonstrating the applicability of the present findings to an explanation of category-specific visual agnosia.

Recognition disorders following acquired brain damage can often affect particular classes of objects and leave others spared. In recent years, one particular form of agnosia that has received much attention is that where the boundary between impaired and intact recognition can be roughly defined along the semantic criterion of biological vs. non-biological objects. In certain patients the recognition of non-biological objects is impaired (Hillis & Caramazza, 1991a, 1991b; Sacchet & Humphreys, 1992; Warrington & McCarthy, 1983, 1987, 1994); in others, the agnosia selectively affects biological items (Basso, Capitani, & Laiacona, 1988; Etcoff, Freeman, & Cave, 1991; Farah, Hammond, Mehta, & Ratcliff, 1989; Farah, McMullen, & Meyer, 1991; Hart & Gordon, 1992; Hillis & Caramazza, 1991a; Humphreys, Riddoch & Quinlan, 1988; Mauri, Daum, Sartori, Riesch, & Birbaumer, 1994; McCarthy & Warrington, 1988; Mehta, Newcombe, & De Haan, 1992; Montanes, Goldblum, & Boller, 1995; Pietrini et al., 1988; Ratcliff & Newcombe, 1982; Riddoch & Humphreys, 1987a; Sartori & Job, 1988; Sartori, Job, & Coltheart, 1992; Sheridan & Humphreys, 1993; Silveri & Gainotti, 1988; Silveri, Daniele, Giustolisi, & Gainotti, 1991; Warrington & Shallice, 1984; Young, Newcombe, Hellawell, & De Haan, 1989).<sup>1</sup>

Category-specific agnosias represent a great challenge to theories of visual object recognition and semantic representation. Specifically, a question of major concern is whether, along the functional pathways involved in object recognition, specialized operations are dedicated to the processing of a particular class of objects (biological or non-biological), or whether category specificity emerges from the interaction between a recognition system without category-specific operations and the particular properties of different classes of stimuli? Then, if the former, what are these category-specific operations? If the latter, what is the nature of the interaction between an observer's perceptual/cognitive system and the structure of the world? The problem of category-specific agnosia

---

<sup>1</sup>In agnosia for biological objects, the categories mainly affected are animals, insects, birds, plants, fruits and vegetables, and faces. However, the disorder is often accompanied by a recognition impairment for particular categories of non-biological objects: foods, musical instruments, and makes of car (Damasio, 1990; Farah, 1991). This extension of deficits for biological items to particular classes of non-biological objects is discussed later in the paper.

is of a sufficient importance that, if solved, it should provide us with a fundamental insight about the operations involved in object recognition.

In this paper, we report the study of a visual agnostic patient with acquired bilateral inferior temporal lobe damage. ELM is severely impaired in the visual recognition of biological objects, but most categories of non-biological items are entirely spared. The results obtained from our studies reveal critical constraints on the phenomenon of visual agnosia specific to biological items; they also provide fundamental information about aspects of the functional organization of the visual recognition system and about some of the psychological dimensions along which the global properties of object shapes are represented. Finally, inferences are drawn as to the possible contribution of the inferior temporal cortex in visual object recognition.

### CATEGORY-SPECIFIC VISUAL AGNOSIA

Several cases of category-specific deficits consequent upon acquired brain damage have so far been reported in literature. However, only a minority of these patients can be classified as genuine cases of visual agnosia. Patients with visual agnosia specific to biological objects constitute the domain of interest for the present paper. The degree to which the observations reported here can be relevant for other forms of category specific deficits remains to be determined.

The theoretical framework used for our distinction between visual agnosia and other forms of category-specific impairments rests on a broad definition of the main processing stages assumed by several theories of visual object recognition (e.g. Biederman, 1987; Humphreys & Riddoch, 1987a; Marr, 1982; Palmer, 1977; Pinker, 1984). These stages are taken to occur in the following sequence: (1) construction of an internal description of the stimulus; (2) access to stored structural knowledge of objects; and (3) contact with semantic knowledge of functional/associative properties of objects. Only after this last stage is it assumed that phonological representations can be accessed for object naming (for a review, see Vitkovitch, Humphreys, & Lloyd-Jones, 1993).

Within this context, the deficits that either spare visual object recognition or similarly affect the processing of both visual and verbal materials—thus suggesting a more central impairment—are not included in the present review. A large proportion of the patients with category-specific deficits fall into these classes (Basso et al., 1988; Hillis & Caramazza, 1991a, 1991b; McCarthy & Warrington, 1988; Pietrini et al., 1988; Sacchet & Humphreys, 1992; Sheridan & Humphreys, 1993; Sirigu, Duhamel, & Poncet, 1991; Warrington & McCarthy, 1983, 1987; Warrington & Shallice, 1984). Other patients have shown positive evidence for an impairment prior to semantic access from vision, which would be congruent with visual agnosia, but their disorder is complicated by a deficit in the verbal domain (Mehta et al., 1992; Ratcliff & Newcombe, 1982; Silveri & Gainotti, 1988; Young et al., 1989). It is worth noting that, in

other cases with a category-specific impairment, the deficit spares object recognition; rather, it specifically affects the retrieval of object names (Farah & Wallace, 1992; Hart, Berndt, & Caramazza, 1985; Pietrini et al., 1988; Sartori et al., 1992). Finally, there are patients who, given the information available, cannot be classified with certainty (Farah et al., 1991; Montanes et al., 1995; Silveri et al., 1991; Warrington & Shallice, 1984).

Excluding the patient studied here, ELM, six unambiguous cases of category-specific visual agnosia (henceforth, CSVA) have been reported in the neuropsychological literature: JB (Humphreys et al., 1988; Riddoch & Humphreys, 1987a); Michelangelo (Sartori & Job, 1988; Sartori et al., 1992); LH (Etcoff et al., 1991; Farah et al., 1989, 1991); KR (Hart & Gordon, 1992); Helga (Mauri et al., 1994) and DRS (Warrington & McCarthy, 1994). With the exception of JB, who suffered a left parieto-occipital lesion, all these patients show temporal or temporo-occipital damage, either bilateral or to the left hemisphere—with, in several cases, more extensive damage elsewhere. In none of these cases is there evidence for a deficit in the encoding of visual stimulus properties that could plausibly account for the agnosic symptoms.

Whereas DRS is mainly impaired for the visual recognition of man-made objects, JB, Michelangelo, LH, KR, and Helga show a severe deficit with biological objects, whereas their recognition of artefacts is largely spared. It appears that a central distinction between CSVA for man-made vs. biological objects can be made on the capacity of patients to access stored structural knowledge from vision and from language. Access to structural knowledge from vision is typically assessed by a reality decision task, where subjects have to discriminate between visual depictions of real vs. unreal (invented) but plausible objects. Access from language is usually assessed either by object recognition from a verbal *description* or by probe questions about the structural properties of an object referred to by its name.

On tasks assessing access to stored structural knowledge, DRS has been shown to perform within normal limits. Thus, he shows no impairment in performing reality decisions or in recognizing objects from a verbal description. In marked contrast, all patients with CSVA for biological objects are impaired, in one way or another, at the retrieval of structural knowledge about objects belonging to categories for which they are agnostic. Thus, in reality decision tasks, LH, Michelangelo, and Helga failed to discriminate between real and unreal biological objects but showed much improved performance with artefacts. However, JB and KR showed normal reality decision performance with biological items. In verbal tasks probing knowledge of biological objects, all five patients exhibited a clear deficit in the retrieval of perceptual properties. In contrast, all of them showed normal or at least relatively preserved performance at retrieving non-perceptual (i.e. functional/associative) knowledge.

As LH, Michelangelo, and Helga failed to retrieve structural properties of biological objects from both pictorial and verbal inputs, their agnosia was attrib-

uted to a damaged structural representation system (Etcoff et al., 1991; Mauri et al., 1994; Sartori & Job, 1988; Sartori et al., 1992). In contrast, JB and KR could retrieve structural knowledge about biological objects from pictures but not from names, and they could retrieve functional/associative knowledge about these objects from names, but obviously not from pictures, because they are agnostic. From this, it appears that the functional impairment in JB (Riddoch & Humphreys, 1987a) and KR affects the pathway responsible for the interactions between the structural and semantic representation systems.<sup>2</sup>

What are the implications of these cases of CSVA for biological objects for the knowledge systems that contribute to visual recognition? Several rival hypotheses have been proposed about the operation of the normal visual recognition system to account for the occurrence of CSVA when this system is damaged by a brain lesion.

Sartori and Job (1988), Sartori et al. (1992), and Etcoff et al. (1991) have proposed that the structural knowledge system may be divided into two or more separate sub-components, each representing structural properties of particular classes of objects. According to the account, LH, Michelangelo, and Helga would have suffered damage to the components representing structural knowledge for biological objects. The speculation proposed by Etcoff et al. (1991) to justify this semantically determined modularity within the structural knowledge system is that the representation of particular categories of objects (e.g. animals) may require specific kinds of structural features, which would not be necessary for the representation of other object classes.

Warrington and her collaborators (Warrington & McCarthy, 1987, 1994; Warrington & Shallice, 1984) have taken a different approach; they attempt to explain category-specific deficits in general rather than only CSVA for biological objects. According to these authors, category-specific impairments are to be attributed to damage to “procedures required for computing meaning” (Warrington & McCarthy, 1994, p. 1472)—that is, semantics. In the case of a visual recognition impairment, a visual semantic system, to be distinguished from stored structural knowledge, would be damaged. This visual semantic system is assumed to be divided according to sub-domains representing either sensory or functional properties.<sup>3</sup> Warrington and her colleagues further propose

---

<sup>2</sup>Although Hart and Gordon (1992) conceded that the explanation proposed here could fully account for their observations, they preferred to argue for a less parsimonious account resting on the assumption of duplicate visual-perceptual representations within structural and semantic knowledge systems. They proposed that brain damage in KR affected the representations of the visual attributes of biological objects that were stored within the semantic system.

<sup>3</sup>A parallel organization is proposed by the authors for a verbal semantic system dedicated to the recognition of verbal materials. Furthermore, a finer grain of analysis of the semantic system has been proposed by Warrington and McCarthy (1987; see also McCarthy & Warrington, 1988) to account for striking dissociations observed in some patients within the broad classes of biological or man-made objects. These further proposals are not represented here as they do not appear central to the present discussion.

that most biological objects are mainly characterized by their sensory attributes, whereas most artefacts are mainly defined by their function. Therefore, semantic damage affecting the representation of sensory attributes would lead to agnosia for biological objects, whereas semantic damage affecting functional representations would be responsible for deficits specific to artefacts. Farah and McClelland (1991) have provided empirical evidence supporting the disparate weightings of sensory and functional properties for the semantic representation of biological and man-made objects. When this disparity is implemented within a connectionist semantic system mediating picture/name associations, Farah and McClelland (1991) have shown that damage to “sensory units” leads to a greater deficit for biological than man-made objects, whereas the opposite effect occurs with damage to “functional units”.

The theory of Warrington and her collaborators seems to account well for the deficits shown by CSVA patients such as KR, JB, and DRS, whose structural knowledge system as well as its access from vision appear to be intact (cf. reality decisions). Indeed, in such cases, it seems one is forced to conclude that the impairment lies either in the access to semantics or within the semantic system itself. However, the theory appears less successful in patients like LH, Michelangelo, or Helga, who fail to perform reality decisions for biological objects but succeed with artefacts. Indeed, failure in the reality decision task seems attributable to an impairment of the system storing structural knowledge, and the patients' performance in this task parallels their visual recognition impairment (e.g. naming). From this, it appears that the functional damage suffered by LH, Michelangelo, and Helga would not be within visual semantics but, rather, at the level of structural knowledge representations.

In contrast to the previous accounts, that proposed by Riddoch and Humphreys for CSVA for biological objects (1987a; Humphreys and Riddoch, 1987b; Humphreys et al., 1988) makes no claims about category-specific operations for visual object recognition or about a division of semantic memory into different classes of knowledge—that is, they proposed unitary structural knowledge and semantic memory systems. They argued that CSVA for biological objects results from a combination of incidental properties of the visual recognition system and of the nature of structural contrasts that exist between semantically related biological items. On the one hand, these authors (Humphreys et al., 1988; for relevant evidence see also Vitkovitch & Humphreys, 1991; Vitkovitch et al., 1993) have shown that different stages involved in object recognition communicate by a cascade process, so that interaction effects occur in normal observers between experimental variables—within category visual similarity and name frequency—thought to influence separate stages of processing (for properties of cascade networks see McClelland, 1979). On the other hand, ratings by normal observers indicated that the perceived structural similarity between semantically related items is greater within biological than man-made categories. It seems, then, that most biological objects would possess

structurally similar neighbours that are also semantically related. In contrast, the sample of man-made objects studied by Humphreys et al. (1988) rarely appeared to have structural neighbours that were also semantically related.

Within the context of a processing system operating in cascade, structural similarity and semantic proximity would be expected to interact non-linearly in visual recognition tasks, such that a combination of structural and semantic proximity—which seems to exist mainly for biological objects—should be much more difficult to resolve than either in isolation. Clearly, a normal recognition system must be able to resolve this conjoined proximity problem. Still, evidence has been reported that, in recognition tasks, greater within-category structural similarity leads to more errors—with brief low-contrast exposures—or longer response times in normal observers (Gaffan & Heywood, 1993; Humphreys et al., 1988). From this, it seems that in a damaged system the most fragile semantic categories should be those comprising exemplars that are structurally close. Riddoch and Humphreys (1987a; Humphreys et al., 1988) tested this hypothesis in JB, a patient with CSVA, using a word/picture matching task in which the subject was given the auditory name of an object, which he had then to select from a set of pictures. JB made frequent errors when distractor pictures were both visually and semantically similar to the target but was normal if, relative to the target, distractors were either visually and semantically dissimilar, or visually dissimilar but semantically related. These results support the conjecture that visual recognition problems involving a combination of structural and semantic proximities are more vulnerable to the effects of brain damage. As noted earlier, it seems that this is exactly the kind of problem that needs to be solved for the visual recognition of biological items.

Critical questions about the account provided by Riddoch and Humphreys for CSVA for biological objects still remain, however. First, the definition of perceptual similarity used by these authors was mainly empirical—that is, judgments by normal observers. Thus, the precise perceptual properties that determine apparent similarity are yet to be identified. Moreover, the procedures by which a visual object or a name can access structural knowledge about objects, as well as the content and organization of these representations, remain to be established. Second, it seems that clarifications need to be brought for the explanation of the effect of within-category perceptual similarity in CSVA for biological objects. Indeed, this variable is characterized by a conjunction of two kinds of codes—one semantic, the other structural—and it seems that each is represented in different systems (see earlier). In CSVA cases where the evidence suggests that it is the pathway between structural and semantic knowledge systems that is affected—i.e. JB and KR—the explanation of the interactive effect of structural and semantic similarities is relatively straightforward within the context of a system operating in cascade, as Riddoch and Humphreys have shown. The situation becomes more complex, however, in CSVA patients with apparent damage to the structural knowledge system itself—i.e. LB,

Michelangelo, and Helga. The fact that these patients cannot reliably distinguish between real and unreal visual instances of biological objects suggests that, even from a visual input, they cannot retrieve stored information about the structure of known objects. Under the assumption of a semantic access that *follows*—that is, is not parallel to—access to structural knowledge, can we expect semantic proximity between items to affect the retrieval of the stored structural representations?

The present report is organized in three main parts. First, we present background information on the patient that is the focus of the study, ELM, as well as the results of a series of experiments—described in detail elsewhere (Decter, Bub, & Chertkow, submitted)—that allow a broad characterization of the functional locus of the impairment he suffers. Second, an exploratory study designed to establish the main constraints that operate on the patient's agnosia is reported. An hypothesis about the precise operation ELM fails to perform in his attempts to recognize biological objects is derived from this study. The experiments constituting the third and main part of this report are designed as tests of this hypothesis. To anticipate on the results, the findings obtained in ELM indicate a deficit specific to the retrieval of structural knowledge about visual objects, which applies even to semantically neutral stimuli. More precisely, the patient is impaired at retrieving the full set of features that are required to specify uniquely the shape of a particular object relative to others from which it has to be distinguished. The results of the final experiment of this series suggest that the category specificity of ELM's agnosia may be explained by the beneficial effect of semantic distance between the objects to be recognized on the retrieval of structural knowledge.

## CASE DESCRIPTION AND FUNCTIONAL LOCALIZATION OF THE IMPAIRMENT

ELM, born in 1928, is a retired dispatcher. His first neurological symptoms were of sudden onset and occurred in December 1982, after hospitalization for atrial fibrillation. At that time, the patient showed nominal dysphasia, left/right confusion, dyscalculia, and agraphia, with no obvious sign of alexia. An emergency CT scan revealed a hypodensity deep in the right mesiotemporal lobe. At discharge, ELM exhibited residual signs of nominal dysphasia and a mild memory impairment, which later disappeared. In August 1985, ELM was readmitted at the Montreal Neurological Institute and was showing a pronounced dysnomia, memory impairment, and dysgraphia. A CT scan conducted after that episode revealed the bilateral inferior temporal lesions shown in Fig. 1. The main complaints of the patient concerned the visual recognition of faces, fruits and vegetables, and colours. As a preliminary indication of the severity of the patient's agnosia, he reported that he could not even recognize his own face in the mirror, or the faces of close relatives he saw on a daily basis.






FIG. 1. Localization of lesions in ELM.

On the WAIS-R, ELM obtained a score in the normal range, with no verbal/performance discrepancy (verbal IQ = 93; performance IQ = 91). His performance on the Wechsler Memory Scale was also within the normal range (MQ = 92), but he obtained a score of 1 in the delayed recall of geometric figures, which is very poor. ELM also showed clear signs of surface alexia and surface agraphia. His reading is characterized by numerous errors in naming irregular words, especially those of low frequency (42% errors). He is also impaired, however, in reading regular words and non-words, although he is capable of identifying the letters of items he misreads. His spelling deficit is manifested by incorrect although phonologically plausible productions (e.g. "bone" spelled as *b-o-a-n*), which occur especially with low-frequency irregular items.

The details of the preliminary experiments designed to investigate the patient's agnosia, which are reported here, as well as of related others, are available in Decter et al. (submitted).

On a variety of visual-perceptual tests, ELM showed a normal performance. In the identification of overlapping line drawings representing man-made objects, the patient made no error (18/18 correct). His matching of pictures of objects taken from different views was also excellent (animals: 7/7 correct; man-made objects: 18/19 correct), as was his naming of pictures of common household objects taken from non-canonical views (ELM: 24/27 correct; average of 10 matched controls: 24/27 correct). In the identification of hierarchical stimuli (large letters made of smaller letters), ELM's performance shows normal effects of global precedence (faster identification of global than local items) and of global-to-local interference (cost associated to a mismatch of local and global information when local items have to be identified). Such observations argue for a normal capacity of generating an adequate internal description of visual input

in ELM. Further, the patient did not show, at any time, signs of visual field deficits. It may be noted, in addition, that the patient drives a car expertly. The experiments described below also provide further indication that the patient's deficit does not lie with the encoding of visual stimulus properties.

Naming of line drawings from the Snodgrass and Vanderwart (1980) set confirmed the patient's complaints about visual recognition and revealed a dramatic performance dissociation between common biological and non-biological objects. Overall, naming of 66 pictures of animals, birds, insects, fruits, and vegetables revealed an error rate of 61%. The biological category on which ELM performed best was that of animals, for which his error rate was of 44%. In contrast, his error rate in naming 79 pictures of non-biological objects (articles of clothing, kitchen utensils, musical instruments, indoor objects, sports objects, and tools) was 12%. Post-hoc analyses established that this dissociation between biological and non-biological objects could not be accounted for by visual complexity or familiarity.

ELM's disorder is not just one of retrieving the name or the semantic information associated with visually presented biological objects. When asked to discriminate between real and unreal but plausible visual exemplars of animals (e.g. a deer head on a cow body), ELM performed at chance (error rate of 41%). In contrast, discrimination between real and unreal versions of man-made objects (e.g. a typewriter with a piano keyboard) was excellent (error rate of 7%). ELM's failure to perform the most basic form of perceptual categorization on biological objects suggests that the patient is impaired at retrieving stored structural information about these items from a visual input.

The failure at retrieving structural knowledge about biological objects was also confirmed in a task where ELM had to answer two-alternative forced-choice probe questions about common biological objects referred to by their auditory names. With questions having to do with the visual appearances of animals (e.g. "does a camel have horns or no horns?"), the patient performed at chance (error rate 45%). In contrast, his accuracy was much better with questions probing encyclopaedic/semantic knowledge about the same items (e.g. "does a camel live in the jungle or the desert?"). His error rate (15%) with these latter questions was slightly higher than that shown by normal controls, however.

There is reason to believe that ELM's moderately impaired performance on encyclopaedic probe questions to animals may not be an invariable aspect of the functional lesion responsible for his poor ability to identify visual exemplars of biological kinds. The dissociation between perceptual attributes and more general functional encyclopaedic knowledge can be seen most dramatically when ELM is tested on fruits and vegetables. Thus, it is clear that he is very seriously impaired in his ability to identify perceptually different members of this category (e.g. tomato, apple, cucumber, etc.). He makes numerous errors on labelling simple line drawings of these objects (59% errors), and his accuracy is

similarly compromised in a word–picture matching task (see Experiment 1). In contrast, his definitions of fruit and vegetable exemplars from the spoken *word* reveal a detailed and rich understanding of *non-perceptual* attributes. Figure 2 illustrates ELM’s obvious difficulties in his efforts to draw from memory the shape of the exemplars cucumber, green pepper, lemon, peach, pear, and strawberry.

His responses when asked to tell us what he knew of different fruits and vegetables on being presented with the *spoken name*, however, are strikingly apposite. Thus:

*Lemon:* “Grows on trees in warm climates, in the Southern part of the United States. You can squeeze lemon juice on fish, or make drinks (a Tom Collins needs a spoonful of lemon), you can have lemon with ice tea, or make lemon pie.”

*Strawberry:* “They grow close to the ground on shrubs. You can make jam with them, have shortcake with strawberries and whipped cream, you can eat them with ice cream or in fruit salad.”

*Tomato:* “Grows on a bush. You get tomato juice from tomatoes. You can slice them and make a tomato sandwich, you can make tomato soup from them, you can fry tomatoes. Good for salads.”

The same task with black-and-white *line drawings* as items reveals (as would be expected) definitions that are clearly based on visual misidentifications. So:

*Lemon:* “Grows on the ground. You harvest them, peel them, slice them, fry them or make salad with them.”

*Strawberry:* “They are red or green, and you eat them by picking them from the bunch and popping them in your mouth.”

*Tomato:* “They grow on trees. You can eat them plain, make sauce with them, or pies.”

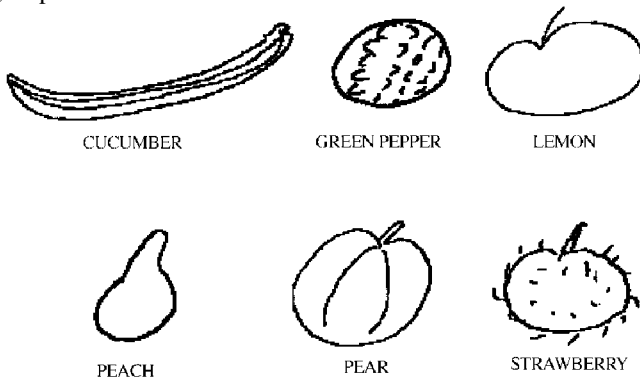


FIG. 2. ELM’s drawing of fruits and vegetables when given the names of the objects.

In the connectionist model of Farah and McClelland (1991) mentioned previously, which has been designed to simulate the phenomenon of category-specific recognition impairments, separate sets of units are assumed to represent functional and perceptual knowledge. Due to the high level of interactivity assumed between these two sets of units, however, damage to the perceptual units only would lead to impaired retrieval of functional knowledge. Our initial observations, together with the experiments that follow, suggest that whatever weak impairment ELM may have in the recovery of functional/associative knowledge, it is not a major determinant of his agnosia.

Overall, the preliminary observations obtained from ELM indicate a visual agnosic disorder that is specific to biological items and cannot be attributed to a deficit in *encoding* the morphological properties of visual stimuli. Furthermore, the patient's performance in the reality decision task and with probe questions on biological items suggests damage to the system representing stored structural knowledge. In these respects, ELM appears as very similar to three other CSVA patients reported in the neuropsychological literature: LH (Etcoff et al., 1991; Farah et al., 1989, 1991), Michaelangelo (Sartori & Job, 1988; Sartori et al., 1992), and Helga (Mauri et al., 1994).

## ERROR PATTERN IN THE RECOGNITION OF BIOLOGICAL OBJECTS

One important stage in the study of the behavioural effects of brain damage is the specification of the error pattern exhibited by a patient. In visual agnosia, possibly the most powerful way to conduct this kind of analysis is by gathering a confusion matrix representing the items the subject finds difficult to distinguish for the purpose of visual identification. It was expected that such data would provide clues about the constraints involved in the recognition of biological objects that ELM fails to resolve successfully.

### EXPERIMENT 1

The stimuli used for this study were line drawings of common fruits or vegetables as well as their auditory names. The task ELM had to perform was one of picture–word matching. Fruits and vegetables were studied instead of other categories of biological objects because they are the simplest items morphologically that we could identify for which ELM is agnosic (59% errors in naming). If, for instance, there is any effect related to the visual properties of objects on confusion errors, this effect would be more apparent and more tractable with structurally simple than with complex objects. Indeed, most fruits and vegetables can roughly be defined as being made of a single component (in the sense of Hoffman & Richards, 1984), whose shape may be described by relatively simple

deformations of a smooth closed contour. In contrast, other categories of biological objects such as animals, insects, plants, and faces are made of multiple components, sometimes with complex shapes. In these cases, multiple determinants related to structural properties may be responsible for identification confusions between objects, and these are less likely to be isolable.

## Method

The stimuli used were line drawings (see Fig. 3) and auditory names of the following 13 fruits and vegetables: apple, aubergine, banana, carrot, cucumber, lemon, onion, orange, pear, pepper, pumpkin, strawberry, and tomato. Visual and auditory stimuli were digitized, and their presentation was controlled by a Macintosh computer equipped with a 14-in RGB monitor. The same equipment served to control the other experiments reported in this paper. Except for aubergine, cucumber, and strawberry, the line drawings were taken from Snodgrass and Vanderwart (1980). Stimulus size (in the standard orientation) on the display screen ranged from 4.0 to 9.0 cm wide and from 2.0 to 6.0 cm tall. The orientation of visual stimuli on the picture plane was varied over 8 possible orientations— $0^\circ$  (i.e. the standard orientation),  $45^\circ$ ,  $90^\circ$ ,  $135^\circ$ ,  $180^\circ$ ,  $225^\circ$ ,  $270^\circ$ , and  $315^\circ$ —randomly and without constraints, from trial to trial. The onset of picture–word pairs on each trial was simultaneous and was preceded by a 1-sec fixation point. The subject's task was to decide whether the stimulus pair referred to the same real world object. The subject was instructed that he could take his time to respond, and the visual stimulus remained visible until the response. The subject was given no feedback as to whether his responses were correct. A total of 1560 trials was run in 10 separate blocks, within which trials were distributed randomly. Matching picture–word pairs were presented on half the trials, and the other half was made of mismatching pairs. The 156 mismatching picture–word pairs possible with the stimulus set used were tested a total of 5 times each, and matching picture–word pairs were tested on a total of 60 trials each. Over the entire experiment, 0.8% of trials were lost due to technical errors.

## Results and Discussion

ELM's error rate for each picture–word pair is presented in Table 1. The analysis of this data set focuses on error rates with mismatching picture–word pairs—false positives—which are more informative than errors with matching pairs. That is, it is quite possible to formulate a hypothesis about the origin of false positives on the basis of the structural or semantic relationships between the items that are confused. The reason why the patient would reject matching word/picture pairs is less obvious. Therefore, we will first discuss the error distribution on negative trials and then examine how the account we are proposing for these may explain errors on positive trials.

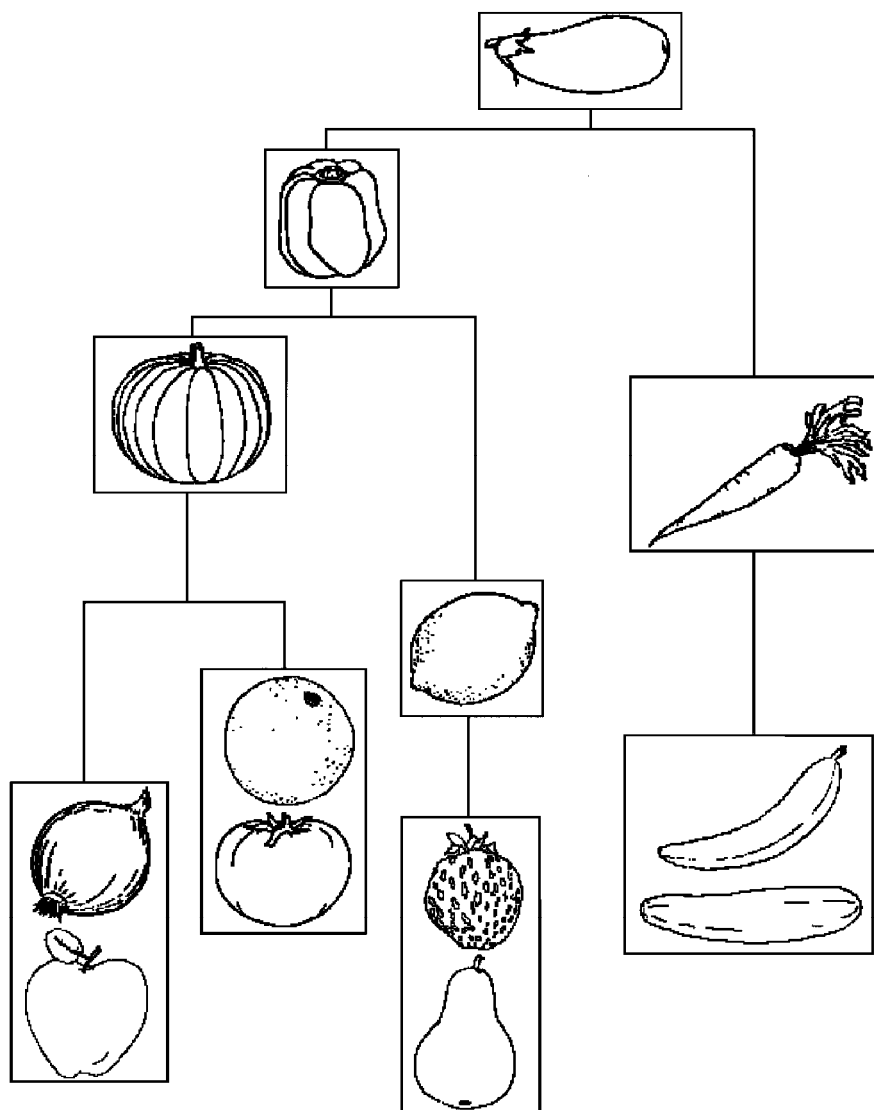


FIG. 3. Line drawings of the fruits and vegetables used in Experiment 1, grouped according to the hierarchical clustering solution obtained on the basis of ELM's performance.

The confusion matrix presents a number of interesting features that suggest that ELM's identification errors are organized around a well-defined principle. The average error rate on mismatching trials is 15.4%. The distribution of these errors across the cells of the confusion matrix is far from random. Indeed, this distribution is greatly uneven across cells, and no error was made on the majority

TABLE 1  
ELM's Error Rates in Matching Line Drawings and Names of Fruits and Vegetables

Word	Picture												
	Apple	Aubergine	Banana	Carrot	Cucumber	Lenon	Onion	Orange	Pear	Pepper	Pumpkin	Strawberry	Tomato
apple	0.06	0.00	0.00	0.00	0.00	0.20	0.20	0.60	0.00	0.00	0.60	0.00	0.80
banana	0.00	0.40	0.02	0.60	0.20	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
carrot	0.00	0.00	0.80	0.01	0.75	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
cucumber	0.00	0.40	0.80	0.60	0.15	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
eggplant	0.00	0.24	0.00	0.00	0.40	0.25	0.00	0.40	1.00	1.00	0.40	0.80	0.20
lemon	0.00	0.00	0.00	0.00	0.00	0.23	0.80	0.20	0.40	0.40	0.20	0.40	0.40
onion	0.20	0.00	0.00	0.20	0.00	0.50	0.13	0.60	0.00	0.20	0.25	0.20	0.60
orange	0.80	0.00	0.00	0.00	0.00	0.00	0.40	0.17	0.00	0.20	0.20	0.20	0.80
pear	0.00	0.00	0.00	0.00	0.20	0.40	0.00	0.00	0.06	0.20	0.00	0.80	0.20
pepper	0.00	0.00	0.00	0.20	0.00	0.00	0.40	0.20	0.20	1.00	1.00	0.60	0.00
pumpkin	0.20	0.00	0.00	0.00	0.00	0.00	0.00	0.80	0.40	0.15	0.20	0.20	0.20
strawberry	0.20	0.00	0.00	0.00	0.00	0.25	0.25	0.00	0.20	0.25	0.00	0.44	0.20
tomato	0.80	0.00	0.00	0.00	0.00	0.00	0.20	0.60	0.00	0.40	0.40	0.40	0.22

(55%) of mismatching picture/word pairs. Thus, 59% of the total number of false positives concentrate on only 15% of the possible mismatching pairs, for which the error rate was 60% or greater.

It is quite apparent that these cells with high error rates correspond to structurally related objects, whereas the semantic relationships between items has no obvious effect. Consider, for instance, the triplet of banana, carrot, and cucumber, which share about the same degree of elongation but differ on other properties. These items were systematically confused with one another. Except for one cell (word = "banana"; picture = "cucumber"), if the name of one of these items was paired with the picture of another, ELM accepted the pair as referring to the same object on 60% of the trials or more. In addition, these items were rarely confused with other objects outside this sub-set. Other comparable indications of an effect of structural similarity on confusion errors, with sub-sets of similar items being frequently confused with one another but rarely with other items outside the sub-set, are present throughout the matrix.

The results of a hierarchical cluster analysis (Fig. 3) support this description of the confusion data. Three main groups of items, which represent sub-sets of objects that were frequently confused, were isolated by this analysis. It is readily apparent that these groups are defined over structural properties. Thus, the group banana/carrot/cucumber appears to be characterized by the fact that all these items are relatively elongated. In contrast, the group apple/onion/orange/tomato is made of items that are relatively round. Finally, it seems that a common property shared by items from the group of lemon/pear/strawberry is that all look as if one or both of their ends had been pinched, thus producing some form of protrusions or bumps on their extremities. Note that onion, although not part of this last group, also presents protrusions on its extremities and was frequently confused with lemon (see Table 1).

What can we learn about ELM's visual agnosia from the pattern of confusion errors he displayed? Most probably, there are several different ways to characterize these observations, and therefore any explanatory hypothesis that may be reached needs to be substantiated by further experiments. What seemed to us the most obvious feature of ELM's errors in picture-word matching is that if the picture displayed only partially matched the shape properties of the object corresponding to the word, he frequently accepted the stimulus pair as referring to the same object. In other words, it appeared that, if the objects referred to by a mismatching picture-word pair shared a critical shape property, ELM tended to ignore the inconsistency between them on other structural aspects and to accept the match. This suggested that ELM's visual identification of fruits and vegetables may rest on partial shape information such that he fails to consider the full range of shape properties that are necessary to specify one particular object uniquely relative to the ensemble of possible items in the task. For example, he accepted the picture of a banana as matching the word "cucumber" on 80% of trials, most probably because these items share about the same degree of



elongation. If the fact that most cucumbers are straight whereas the picture shown was of a bent fruit had been considered, however, this picture–word pair would have been rejected.

From these observations, it was hypothesized that ELM's agnosia for fruits and vegetables may follow from a failure of exhaustive processing of the multiple properties specifying the shapes of objects—we refer to this as a shape integration impairment. The following series of experiments test this hypothesis. First, however, ELM's results on positive trials, as well as two other features of the confusion matrix, should be discussed.

On positive trials, the patient's average error rate was 19.3%. It would appear that these errors may find an explanation under the hypothesis of a shape integration impairment in ELM. Thus, the notion of partial processing of shape features obviously means uncertainty as to the identity of the object presented—more than one known object is compatible with the limited set of features processed. The high error rates shown by ELM with structurally related mismatching picture–word pairs suggest that, among the different object identities compatible with the limited set of shape features processed, he often, though not always, accepted the alternative named by the word presented. We propose that the same failure may have occurred on positive trials. For instance, with the positive picture–word pair *apple/apple*, the processing of a limited set of visual features must have led ELM to consider alternative identities, such as *orange*, *pumpkin*, or *tomato*, as compatible with the picture shown. Although on most occasions the patient accepted the positive picture–word pair, he rejected it on some trials, presumably to the benefit of an alternative structurally related object.

One important point to underline, which can be noted in Table 1, is that ELM's deficit does not seem to spare any of the items tested in Experiment 1. Thus, on each row and each column of the matrix, there is at least one cell that shows a substantial number of errors. Although the quantity of data is limited, what it suggests is that the recognition deficit may not affect item-specific representations but may, rather, involve a relatively general procedure that operates either on internal descriptions of stimuli or on representations of knowledge about object shapes. This impression is confirmed by the work reported further on.

Another interesting feature of the confusion matrix is its asymmetry, which is present for several mismatching pairs of items. Consider, for instance, the pair *strawberry* and *pear*. When the picture of *pear* was presented with the name of *strawberry*, ELM's error rate was only 20%. In contrast, with the picture *strawberry* and the name *pear*, the error rate increased dramatically to 80%. Other pairs of items exhibit an asymmetry of comparable or greater magnitude: *cucumber*/"banana", *aubergine*/"pear", *aubergine*/"pepper", *aubergine*/"strawberry", and *orange*/"pumpkin". This kind of asymmetry in confusion matrices has been reported in the study of visual letter recognition (e.g. Loomis, 1982).

Even though the following experiments do not focus on this question, we would like to underline the analogy between an asymmetry within confusion matrices, such as that reported here, and the asymmetry of performance of normal observers in the visual search for particular kinds of features. For example, Treisman and Gormican (1988) have shown that the search for an oblique bar among vertical lines was performed by a spatially parallel process, whereas a vertical displayed among obliques was searched serially. This type of result has been interpreted as a reflection of the way the visual system codes stimulus properties on some perceptual dimensions (Treisman & Gormican, 1988). In particular, it is assumed that whereas some feature values would be coded as a positive signal (e.g. oblique), others would be registered as the default (i.e. neutral) value on that dimension (e.g. vertical) and would therefore correspond to a non-signal code. It seems possible that a comparable principle may apply to visual object recognition and that it would be responsible for the asymmetries observed in the present confusion matrix. Thus, for instance, it is conceivable that in some cases the *presence* of a particular property in the picture displayed was sufficient to veto the picture–word pair A/B, but that its *absence* with the pair B/A was not explicitly signalled and therefore led ELM to accept the match erroneously.

## SHAPE INTEGRATION IN CATEGORY-SPECIFIC VISUAL AGNOSIA

The main finding from Experiment 1 was that ELM's identification confusions in a picture–word matching task appeared to occur essentially between items that possess some shape properties in common. This led to the hypothesis that his agnosia may be related to a failure in integrating the full range of features that are required to uniquely specify the shape of one particular object relative to others that are structurally related to it—i.e. a shape integration impairment. As noted above, this would explain why, for instance, ELM confused banana and cucumber, which share about the same degree of elongation, in spite of their obvious difference on bending. Most of the experiments reported in the present section were designed as tests of this hypothesis. The last experiment of the series will attempt to determine how the notion of a shape integration impairment may relate to the category specificity of ELM's visual agnosia.

One critical constraint that needs to be met for an *experimental* assessment of shape integration concerns the control over the structural relationship between the items used in a particular task. Indeed, if there is such an operation as the integration of shape properties, it will be involved only in tasks where correct performance requires processing over two or more aspects (i.e. dimensions) of shape. This requirement exists only when the items used have some of their properties in common. In the present case, to allow the formulation of specific predictions, these shared properties need to be established a priori rather than

post hoc, as was the case in Experiment 1. Considering this, the approach we selected was to design a method for the parametric manipulation of shape properties over a number of relevant shape dimensions, which allowed the generation of synthetic stimuli with fruit/vegetable-like shapes.

Mainly from an examination of the most salient shape properties of a variety of fruits and vegetables and some of the confusion errors made by ELM in Experiment 1, we hypothesized three simple dimensions that can serve to define global changes in the shape of a smooth closed contour (i.e. ellipsoid):

- elongation, i.e. ratio of minor to major axes of the ellipsoid.
- global tapering along the major axis of the ellipsoid.
- global curvature (i.e. bending) of the ellipsoid perpendicular to the major axis.

By combining particular values (i.e. features) on each of these dimensions, a wide variety of discriminable candidate shapes that resemble fruit and vegetable forms can be synthesized. Each of these potential stimuli corresponds to a particular point in a three-dimensional shape space defined over the properties of elongation, tapering, and curvature. For the generation of experimental stimuli, these dimensions were defined operationally by the sequential global transformations described in the Appendix.

As examples of the shapes to which these dimensions correspond, consider the shape of an idealized cucumber as a starting point, which we will define as elongated, not tapered, and not bent (i.e. straight). To go from the shape of a cucumber to that of a watermelon, a change in elongation is needed. To go from cucumber to carrot, the initial shape would have to be tapered. Finally, if we were to bend the cucumber shape to a certain degree, we would obtain a stimulus looking very much like a banana.

An interesting and important point to note is that the dimensions of deformation defined here are closely related to the intuitive parametric deformations typically associated with synthetic solid modelling (for example, using Platonic solids or superquadrics—e.g. Barr, 1991; Biederman, 1987; Brooks, 1981; Pentland, 1985, 1988). This argues in favour of the applicability of the set dimensions defined above for the representation and perceptual extraction of shapes. Furthermore, Biederman (1987) has presented arguments for the psychological validity of shape dimensions resembling those used here. Some of the observations reported in these experiments provide empirical support for this latter claim.

In all the following experiments, the stimuli used were filled blobs, which were synthesized by the selection of particular parameters on the dimensions of elongation, tapering, and curvature. In addition, in all experiments, orientation of the items on the picture plane was varied over 8 possible orientations— $0^\circ$  (i.e. vertical major axis),  $45^\circ$ ,  $90^\circ$ ,  $135^\circ$ ,  $180^\circ$ ,  $225^\circ$ ,  $270^\circ$ , and  $315^\circ$ —randomly and

without constraints, from trial to trial. The length of the major axis of stimuli was also standardized<sup>4</sup> in order to prevent local cues from contributing to performance.

## EXPERIMENT 2

As noted previously, there are probably several different explanatory hypotheses that may be formulated for the confusion matrix reported in Experiment 1. Thus, our initial impression that confusions essentially occurred between items with common structural features, which led to the hypothesis of impaired shape integration in ELM, requires an experimental verification. This is the goal that led to Experiment 2.

The task used was, as in Experiment 1, picture–auditory-word matching. The main feature of Experiment 2 is that, on negative trials, the structural relationship between the test picture presented and the object corresponding to the word (called the target object) was exactly specified along the shape dimensions described previously. Obviously, this kind of specification could not be achieved in Experiment 1, because the pictures that could be presented on negative trials were constrained by the use of natural objects. In Experiment 2, this constraint existed only for the matching picture–word pairs, where the visual stimulus had to correspond to the form of the object referred to by name. The test pictures for mismatching pairs did not have to correspond to any known object and were determined only by their structural relationship to the target object. These structural relationships on negative trials were of three classes: the test picture differed from the shape corresponding to the word presented along one, two, or three dimensions. Examples of stimuli are given in Fig. 4. As can be determined from this figure, contrasts between the test stimulus and the target object on any dimension where a difference existed were perceptually salient.

From this design, a clear prediction can be made for a subject with impaired shape integration. Indeed, if only a sub-set of shape properties is considered by the subject to perform the matching task, the error rate will be an inverse function of the number of dimensions along which the test picture differs from the target object. Examine, for instance, the stimuli shown in Fig. 4. If the word presented is “banana” and the subject uses only curvature as a criterion to decide whether the picture–word pair matches, the error rates on negative trials should be 66%, 33%, and 0% with 1D, 2D, and 3D difference between the test picture and the target object, respectively.

---

<sup>4</sup>Initially we considered standardizing the surface area of stimuli instead of their length. This, however, proved greatly problematic in experiments where large differences of elongation existed between items. Indeed, to compensate for the loss of surface area of elongated (i.e. thin) items, they had to be extremely long. Apart from the fact that some of these items were so long that they did not fit on the display screen, it also appeared that stimulus length was more likely to serve as a local cue than surface area, given its greater salience.

## Methods

The task was the same as in Experiment 1 (picture–auditory word matching), with simultaneous onset of the auditory word and the test picture. The auditory target words used were “banana”, “carrot”, “cucumber”, and “eggplant” {aubergine}. All test pictures were generated from the three-dimensional shape scheme described previously (dimensions of elongation, curvature, and tapering). All stimuli in this and the next experiments had the length of their major axis normalized at 7.0 cm. On positive trials, the test picture closely matched the shapes of the objects referred to by the target word and were found to be recognizable as such by normal observers. On negative trials, the test picture (see Fig. 4 for examples) differed from the target object on one (elongation, curvature, or tapering), two (elongation–curvature, elongation–tapering, or curvature–tapering), or three (elongation–curvature–tapering) shape dimensions. These shape differences were of a categorical nature (e.g. thin vs. fat; curved vs. straight; tapered vs. not tapered) and were perceptually salient. A total of 504 trials was run in 4 experimental blocks, with half the trials being positive. The order of trials was random, and all conditions were tested in each block. Each type of difference between the test picture and the target object described above was tested on a total of 36 trials. Each trial was preceded by a 500-msec fixation point. The subject could take his time to respond and had to indicate





Condition	Examples of test-picture(s)
Positive trial	
3D difference	
2D difference	
1D difference	

FIG. 4. Examples of the stimuli used in the different conditions of Experiment 2 when the target word was “banana”.

whether the auditory name and the visual stimulus referred to the same real world object. The visual stimulus remained visible until the subject's response. No feedback was given as to the accuracy of responses.

## Results and Discussion

The average error rates shown by ELM in each condition are presented in Table 2. As was the case in Experiment 1, the data analysis focuses on negative trials, because, as noted previously, a clear interpretation is difficult to derive from errors with matching picture–word pairs.

As predicted from the hypothesis of impaired shape integration in ELM, the results of Experiment 2 indicate a major effect of the number of dimensions along which the test picture differed from the target object with mismatching picture word pairs,  $\chi^2(2) = 18.5$ ,  $p < .001$ . Furthermore, the effect is in the predicted direction. Thus, when the test picture differed from the target object on only one shape dimension, ELM accepted the match on almost half the trials—his performance in this condition does not differ from chance,  $\chi^2(1) = 1.8$ ; n.s. As the number of dimensions by which the test picture differed from the target object increased, ELM's error rate decreased accordingly.

As noted in the discussion of Experiment 1, it appears that ELM's errors on positive trials may be explained by the hypothesis of a shape integration impairment. In addition, we should point out that the error rate on positive trials in Experiment 2 (23.0%) is only marginally larger than that observed in Experiment 1 (19.3%). This substantiates our assumption that the positive synthetic shapes used in the present experiment were as acceptable depictions of the target objects as were the line drawings that served in Experiment 1.

Another important point to note in ELM's data from Experiment 2 is the large variability in his performance on negative trials according to the dimension by which the test picture differed from the target object (Table 3). Thus, with a 1D difference, performance was severely impaired with a test picture differing from

TABLE 2  
Error Rate on Negative Picture–Word Matching Trials  
as a Function of the Number of Dimension(s) by  
Which the Test Picture Differed from the Target Object

<i>Condition</i>	<i>%</i>
positive trials	23.0
3D difference	8.3
2D difference	25.0
1D difference	43.5

the target object on global curvature, close to chance with a difference on tapering, and much improved with a difference on elongation. This effect of the dimension by which picture/word pairs mismatched was highly significant  $\chi^2(2) = 41.7, p < .001$ . A congruent pattern of performance was observed with a 2D difference between the test picture and the target object,  $\chi^2(2) = 24.9, p < .001$ . These observations suggest the existence of some hierarchy or biasing in ELM's processing of the dimensions used here to define the shapes of the stimuli. Thus, he seems to have processed information about elongation rather consistently in his picture–word matching, so that when the test picture differed from the target object on that dimension, he noted the difference on almost all occasions and responded accordingly by rejecting the picture–word pair. In contrast, it appears that differences on global curvature were almost never noticed by ELM, so that he accepted the great majority of these picture–word pairs.

To summarize, Experiment 2 provided two main results: (1) ELM's error rate decreased gradually with increases in the number of dimensions by which the test picture differed from the target object. As explained earlier, this observation supports the hypothesis that the patient's agnosia is related to a failure in processing the full range of features necessary to specify the shape of a particular object uniquely relative to others that share some of its properties. (2) ELM's error rate varied greatly as a function of the dimension by which the shape of the test picture differed from the target object. This result indicates that the patient is much more likely to omit processing some properties of the objects for recognition (global curvature) than others (elongation). This last point raises the possibility that the phenomenon to which we refer as a shape integration deficit in ELM—i.e. failure to consider all the properties required for the unique specification of the item—may, in fact, not be directly related to the failure of an integration process, as initially assumed. Rather, it could simply be due to the patient's incapacity for processing information along some of the dimensions by which the shapes of fruits and vegetables may be characterized. Experiment 3 examines this possibility.

TABLE 3  
Error Rate on Negative Picture–Word Matching Trials  
as a Function of the Particular Dimension(s) by Which  
the Test Picture Differed from the Target Object

<i>1D Difference</i>	<i>%</i>	<i>2D Difference</i>	<i>%</i>
elongation	8.3	elongation–tapering	2.8
tapering	38.8	elongation–curvature	19.4
curvature	83.3	tapering–curvature	52.8

## EXPERIMENT 3

From the results of Experiment 2, the possibility arose that ELM's incapacity to use the full set of features necessary to specify the shape of an object in his recognition attempts may be related to an impairment in the processing of some of the dimensions by which fruit and vegetable forms can be characterized. In particular, the evidence suggested an almost total incapacity to process information about global curvature and a severe deficit for the dimension of tapering. The most straightforward prediction made by this hypothesis is that, if ELM was asked to perform object classifications based on a single shape dimension, his performance should be very poor with stimuli differing only on global curvature or tapering. However, he should do notably better with classifications based on elongation. Experiment 3 is a direct test of this prediction.

Each stimulus set used represented a continuous change on either elongation, curvature, or tapering between two poles which correspond to the shapes of known fruits or vegetables (Fig. 5). For instance, on the dimension of elongation, one extreme of the continuum corresponded to the shape of a cucumber, and the other to the shape of an asparagus. On each trial, the subject was required to determine the category (e.g. for elongation, asparagus vs. cucumber) to which the individually presented item belonged. Within this design, if ELM is incapable of encoding information about a particular dimension or cannot use this information for categorization, his classifications should have a random correspondence with variations in stimulus qualities.

### Method

Three sets of 12 stimuli each representing continuous changes on the dimensions of elongation, curvature, or tapering were generated (see Fig. 5). For each stimulus set, the shapes corresponding to the extremes of the continuum were recognizable depictions of a fruit or vegetable (elongation: cucumber/asparagus; curvature: cucumber/banana; tapering: cucumber/carrot). The 10 other stimuli between the extremes of each continuum corresponded to equally spaced (as determined by the parameters used by the shape generation scheme—see Appendix) changes on the dimension of interest. Shape sets were tested separately, each in 4 blocks of 12 trials. On each trial, one of the shapes was shown by itself at the centre of the display screen, and the subject was asked to determine which end of the tested continuum it resembled most. For instance, with stimuli from the elongation set, ELM had to decide whether the picture resembled more a cucumber or an asparagus. The order of trials in each block was random, with each shape of the set tested appearing once. The subject received no feedback as to the accuracy of his responses and was not informed of the dimension tested.






FIG. 5. Stimuli used in Experiment 3: (A) elongation set; (B) curvature set; (C) tapering set.

## Results and Discussion

ELM's classification results are shown in Fig. 6. It is quite evident from these graphs that the patient's responses were systematically related to stimulus properties on each of the dimensions tested. This is confirmed by statistical tests on the distribution of responses as a function of the properties of the shapes presented:  $\chi^2(11) = 29.0$ ,  $p < .01$ , for elongation and curvature;  $\chi^2(11) = 36.0$ ,  $p < .001$ , for tapering. For each stimulus set, one extreme of the continuum consistently received one of the available labels, and the other extreme of the continuum received the other label, with a relatively sharp shift in the categorization at the middle of the range.

The systematic relationship between ELM's classifications and the stimulus properties in Experiment 3 clearly indicate that he is capable of perceptually encoding information about the elongation, tapering, and curvature of visual shapes. The results also indicate that ELM can use these shape dimensions for the categorization of visual items and therefore that he retains knowledge about the critical shape properties of the objects that were tested in each set (e.g. that a banana is curved). From this, the possibility that ELM is incapable of processing some of the shape dimensions studied here, which was suggested by particular aspects of the results of Experiment 2, can be rejected. The issue of ELM's failure to use information on tapering and curvature in the previous experiment is addressed in the discussion of Experiment 4 and in the General Discussion.

It is worth noting that the present observations are in opposition to the speculation proposed by Etcoff et al. (1991) to account for category-specific visual agnosia. These authors hold that specific kinds of structural features, which are

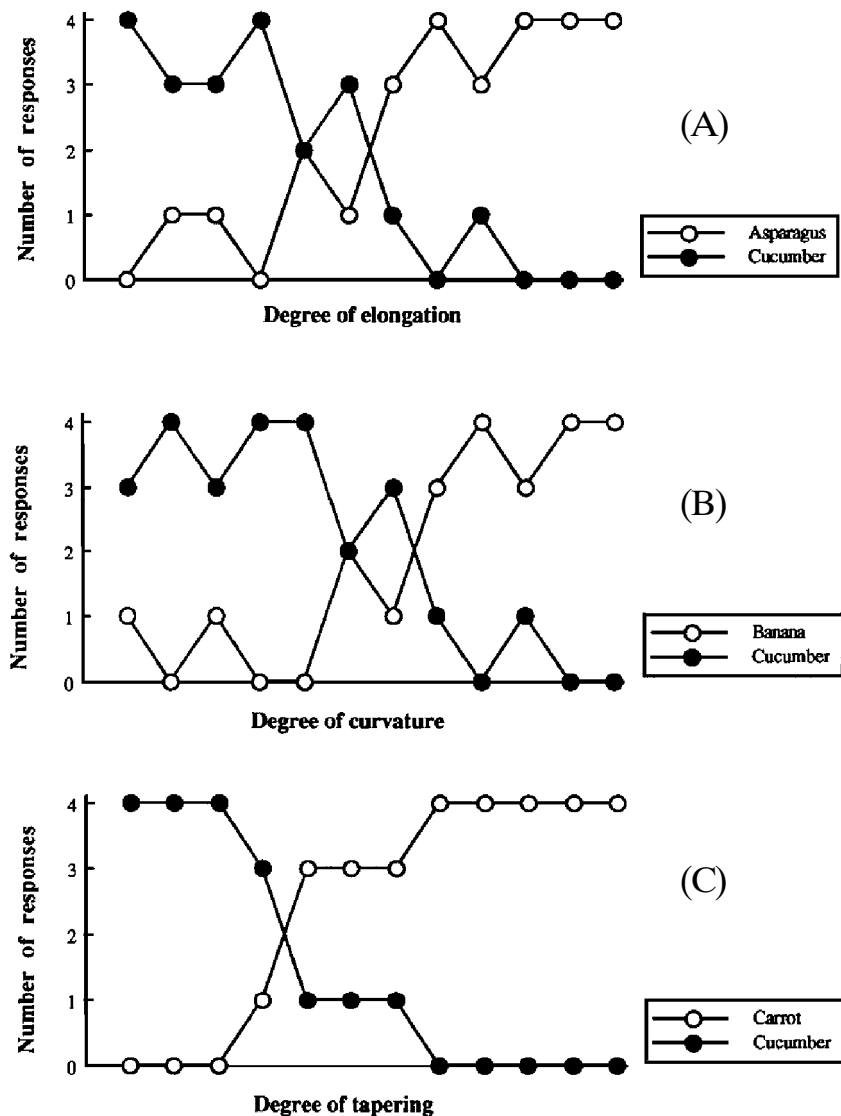


FIG. 6. ELM's classification performance as a function of stimulus properties in Experiment 3: (A) elongated set; (B) curvature set; (C) tapering set.

necessary for the representation of biological objects, are no longer available within the structural representation system of patients with category-specific agnosia. The accurate categorization performance of ELM with shapes varying along the dimensions of elongation, tapering, and curvature suggests that, as far as these dimensions are concerned, the hypothesis of Etcoff et al. does not hold.

It is also interesting to note that, from the observations of the present experiment alone, one would not have inferred that ELM suffers from a visual agnosia for fruits and vegetables. Indeed, these results are in striking contrast to those of the picture–word matching task reported above (Experiment 1), where the cucumber/banana and cucumber/carrot pairs were very often confused. These pairs certainly did not lead to the same degree of confusion in the present experiment. One main feature of the categorization task used here is that adequate performance only required the processing of one of the dimensions of the shapes presented—that is, elongation, tapering, or curvature. It appears that this was not the case in Experiment 1, as the unique characterization of the shape of a particular object relative to the other items used in that task seemed to require the processing of two or more shape dimensions. In other words, consideration of stimulus values on only one shape dimension appeared insufficient to specify uniquely any of the 13 objects tested in that experiment. In fact, the analysis of the confusion matrix obtained from Experiment 1 indicated that the bulk of ELM's confusion was between objects that had at least one particular shape feature *in common*. It is precisely this aspect of the confusion matrix that led to the hypothesis of a failure of shape integration in ELM. The picture–word matching task used in Experiment 2 had been designed as a specific test of this hypothesis and provided support for it. Thus, error rates on negative trials were inversely related to the number of shape dimensions by which the picture presented differed from the target object referred to by the word.

From these observations it appears that, if indeed ELM suffers from a shape integration impairment, one critical constraint on his visual categorization accuracy should relate to the dimensionality of the stimulus set tested. This effect of dimensionality is examined in Experiment 4.

## EXPERIMENT 4

The contrast between ELM's results in Experiments 1 and 2 and in Experiment 3 suggested an effect of the dimensionality of the stimulus set on his capacity for visual categorization, with improved performance when processing of a single shape dimension is sufficient to execute the task successfully. The verification of this possibility by a systematic study of the effect of the dimensionality of the stimulus set would constitute direct support for a failure of shape integration in ELM. Indeed, by definition, the unique specification of a particular object taken from a set made of items differing from one another on only one shape dimension does not require any integration of shape properties. In contrast, integration will be necessary if a stimulus set is made up of variations of two or more shape dimensions and if the items in that set share some of their properties with one another (i.e. if the stimulus set constitutes a conjunction problem).

These two classes of stimulus sets were tested in Experiment 4 within the context of a recognition task where ELM had to learn and then recall the random

assignments of four shapes to four different locations on a computer screen. The selection of this task was motivated by the fact that it corresponds to the most primitive form of visual recognition one can imagine—that is, the assignment of a particular perceptual instance (shape) to a past experience (location). In particular, it should be emphasized that the task involves no semantic or verbal processing and does not refer to any known object. Because of this, any contrast that we may observe between ELM's performance according to the dimensionality of the stimulus set will only be attributable to the visual properties of the shapes tested and not to other non-visual features associated with the items.

Sets of four stimuli each were synthesized from the three-dimensional shape generation scheme described previously. Three stimulus sets corresponded to changes on only one shape dimension (Fig. 7a). Discriminations between items in any of these “single-dimension” sets rest exclusively on one shape dimension and therefore do not require any integration of the different properties of the stimuli. Three other sets were generated by the variations of stimulus properties on two shape dimensions (Fig 7b). These shape sets were constructed such that consideration of the conjunction of the critical shape properties was required to distinguish between items—that is, no single shape dimension could serve in isolation to specify individual items. In other words, the integration of two shape properties was required to distinguish items successfully.

In this experiment, two predictions can be made under the hypothesis that ELM suffers from an impairment of shape integration: (1) His error rate (i.e. pointing to a wrong location) should be notably higher with the conjunction sets

#### A) Single dimension sets

Elongation



Curvature



Tapering



#### B) Conjunction sets

Elongation/Curvature



Elongation/Tapering



Curvature/Tapering



FIG. 7. Stimuli used in each condition of Experiment 4.

than with the single dimension sets, and (2) his errors with conjunction sets should essentially involve confusions between items that share the same feature on one of the critical dimensions (e.g. the target presented is thin and curved and the subject points to the location assigned to the thin and straight item). That is, ELM should rarely, if ever, confuse items that share no critical property.

## Method

Six shape sets of four items each were used. Three of them (single dimension sets) corresponded to equally spaced intervals on one of the dimensions used in the shape generation scheme described previously (elongation, curvature, or tapering). In these single-dimension sets, the properties of the stimuli on the irrelevant dimensions, remained constant across items. In shape space, the structural relationships between items from a “single-dimension” set can be represented as a straight line, which runs parallel to the main axis of the space that corresponds to the dimensions of interest for that set (Fig 8a).

Three other sets (conjunction) were generated by introducing variations on two of the shape dimensions (elongation/curvature, elongation/tapering, curvature/tapering) such that any item had the same value as another on a particular dimension (e.g. elongation), and the same value as a third item on the other relevant dimension (e.g. curvature). The properties of the stimuli on the irrelevant dimension (e.g. tapering) remained constant across all items of a set. Consider, for example, the items used in the elongation/curvature set illustrated in the top row of Fig. 7b. The first item in that row has the same degree of curvature as the second and the same degree of elongation as the third. These structural relationships between shapes exist for all items of the set. In shape space, structural relationships between stimuli in conjunction sets can be presented by a rectangle (or square) with its sides running parallel to the axes of the space that correspond to the dimensions that are varied (Fig. 8). To pursue the example of the elongation/curvature set, it may be noted that the processing of elongation only would specify two different objects, as would that of curvature.

Each stimulus set was tested separately in 10 consecutive blocks of 8 trials each. Before each block, the four stimuli of the set tested were presented simultaneously, each at one corner of the display screen. The shape/location assignments were random and remained fixed throughout the session. The subject was instructed to try to remember the shape/location assignments, and items remained visible as long as he wished. Stimuli were then removed from the screen, and the stimulus set was tested, with each item appearing twice in a block of randomly ordered trials. On each trial, one item appeared at the centre of the screen, and the subject’s task was to point to the location to which it had previously been assigned. The subject could take his time to respond, and no feedback was given about response accuracy.

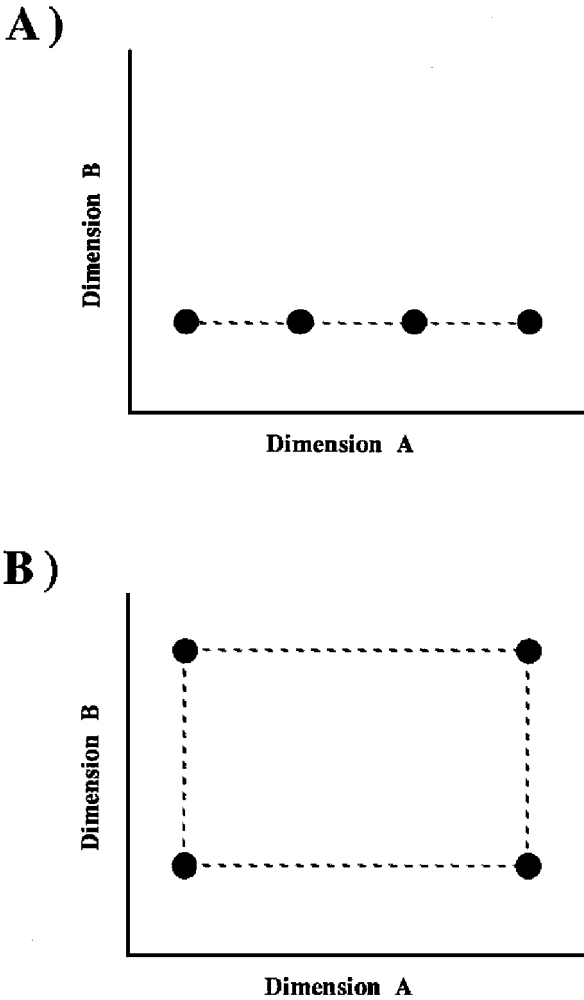


FIG. 8. Structural relationships between items in the single-dimension (A) and conjunction (B) conditions of Experiment 4.

## Results and Discussion

As predicted from the hypothesis of impaired shape integration in ELM, the patient made many more errors with the conjunction sets (56.7%) than with the single-dimension sets (29.2%)—a difference that is highly significant,  $\chi^2(1) = 37.0$ ,  $p < .001$ . This effect of dimensionality cannot be explained by the structural proximity between items in conjunction sets. In fact, as can be seen in Fig. 7, items from single-dimension sets were perceptually much closer to one

another than were items from conjunction sets—this is supported by the parameters used to construct the shapes. If perceptual similarity was the only constraint on performance, one should expect more frequent errors with the single-dimension than the conjunction sets,—an error pattern that is the opposite of that observed in ELM. To demonstrate that the conjunction task is not intrinsically more difficult than that involving a single dimension, Experiment 4 was administered to a new patient, JT, who had suffered a left occipital lesion but who is not agnosic. JT exhibited a greater error rate with the single-dimension (12.9%) than with the conjunction sets (2.5%),  $\chi^2(1) = 18.3, p < .001$ . It should be underlined that JT's pattern of results is congruent with that of normal observers, who, in the perceptual learning of object sets larger than those used here, show enhanced performance with an increase in the dimensionality of the set (Lockhead, 1970). The dissociation seen here between patients as well as the design of stimulus sets exclude an explanation of the dimensionality effect observed in ELM on the basis of perceptual similarity between items. Rather, these facts indicate that ELM's greater error rate with conjunction sets must be attributed to a failure of shape integration.

An analysis of ELM's confusions provides further support for the hypothesis of a shape integration impairment by verifying our prediction about the distribution of errors in the conjunction condition. Thus, with conjunction sets, ELM confused items that had no property in common on the critical dimensions on only 8.8% of the 136 trials on which he committed an error. By chance alone, one-third of the patient's errors would have been of this category. The difference between ELM's error pattern with conjunction sets and a chance distribution of errors is highly significant,  $\chi^2(1) = 36.8, p < .001$ . In other words, the proportion of the errors made by ELM with conjunction sets that involved confusions between items with one of their critical properties in common is greater than would be expected by a chance distribution of errors. It appears that this error distribution cannot be explained by the greater similarity between stimulus pairs sharing one feature than between those sharing no feature. Thus, a detailed examination of ELM's errors indicates that, with the single-dimension sets, the patient never confused items that were as perceptually distant as those sharing one of their features in the conjunction sets.

Another aspect of ELM's performance in the conjunction condition that is worth noting concerns the distribution of confusions between items with one property in common as a function of the particular shape dimension along which they differed. Again, as was noted previously in Experiment 2, confusions rarely occurred between items differing on elongation. Thus, with the elongation/curvature set, confusions between items differing only on elongation occurred only on 6 trials, whereas confusions on curvature were observed on 40 trials. Similarly, with the elongation/tapering set, 7 errors involved elongation, and 34 errors involved tapering. Interestingly enough, however, the results again cannot be taken as an indication that ELM is simply incapable of processing

information on the dimensions of curvature and tapering. Thus, with the curvature/tapering set, the patient confused items differing on curvature on only 3 trials, whereas confusions involving tapering occurred on 33 trials. Overall, these observations seem congruent with other indications of a shape integration failure in ELM. That is, being incapable of using information on both critical dimensions to perform the conjunction task, the patient seems to give one a higher priority and will therefore make few errors on that dimension. When the task involves conjunctions of information on elongation and another dimension, ELM appears to focus systematically on elongation. However, in the curvature/tapering task, curvature is allocated a high priority, and very few errors involve that dimension. This is in marked contrast to the elongation/curvature task, where the patient appeared entirely incapable of making any use of information on curvature.

The results of Experiment 4 have provided strong evidence for a shape integration impairment in ELM. First, the patient's error rate in a visual recognition task was dramatically increased when the unique specification of a particular item in the stimulus set required the integration of information along two shape dimensions—that is, conjunction sets. Second, the errors exhibited by ELM with conjunction sets are of the type one would expect given a shape integration deficit, with confusions occurring mainly between shapes sharing one of their critical properties. A detailed analysis of the patient's performance in Experiment 4 as well as the testing of an additional non-agnosic brain-damaged patient provided strong arguments against an explanation of these two results in terms of the overall perceptual similarity of the shapes that were confused. Thus, the perceptual similarity hypothesis would have predicted a greater error rate with single-dimension sets, as found in JT, and in direct opposition to the observations from ELM. Furthermore, it was observed that confusions between items of conjunction sets involved shapes that were perceptually very distinct; items with a comparable degree of perceptual difference were never confused by the patient within the context of single-dimension sets. These observations imply that the effects observed in ELM are not simply the function of a loss of resolution in shape processing but, instead, reflect the failure to process, or take into consideration, information along multiple shape dimensions for object recognition.

It should also be added that the deficit revealed in ELM by Experiment 4 is one that essentially concerns the processing of shapes and in no way involves components of semantic or verbal processing. Indeed, the task used in the present experiment was entirely non-verbal and involved exclusively the episodic learning of specific associations between shapes and locations. In addition, no semantic content was given to the shapes used, which therefore did not refer to any prior knowledge the patient had. One clear implication of this is that if a shape integration impairment is to be related in any way to ELM's visual agnosia for biological objects—a question addressed in Experiment 6—it must



be assumed that one significant component of this visual recognition disorder lies entirely within the domain of visual shape processing.

The observation of a shape integration impairment in ELM has two critical and related implications, which concern the normal functional architecture of the visual recognition system. Indeed, the observation that brain damage can affect the integration of the properties that characterize objects along several different shape dimensions indicates that such a shape integration process does occur in normal object recognition—cf. transparency assumption (e.g. Caramazza & McCloskey, 1988). By extension, this also means that normal object recognition cannot be conceived as the simple mapping of a holistic description of the visual input onto an internally stored image-like representation of structural knowledge about the object. Rather, our observations show that, between input and access to the appropriate stored representation, an assembly process needs to take place, whereby the discrete features of that object along various shape dimensions are conjoined.

What remains unknown, however, is at which processing stage(s) the shape integration operation takes place. Indeed, it is conceivable that it may concern the the processing of the actual shape of the stimulus an observer is trying to recognize. That is, shape integration may be required for the construction of an adequate structural description of the input. This hypothesis would be analogous to Treisman's feature integration theory (Treisman & Gelade, 1980; Treisman & Sato, 1990), which postulates that the adequate *perception* of feature conjunctions requires that separately encoded stimulus properties be actively conjoined. Alternatively, shape integration might concern the retrieval of stored structural knowledge about objects. Under this hypothesis, the various features that represent our knowledge of the shape properties of objects would be stored separately, and these would need to be conjoined for adequate visual recognition. Obviously, from the study of ELM, it is not possible to exclude either of these hypotheses. What can be done, however, is to determine at which stage (i.e. shape perception or retrieval of shape knowledge) shape integration does occur by specifying the level of processing at which the patient's impairment is manifest. This is the goal of Experiment 5.

Before moving on to the next experiment, one additional and interesting implication of the results of Experiment 4 should be mentioned. Thus, in the present experiment, a clear effect of dimensionality of the stimulus set was observed in ELM with shape parameters defined along the dimensions of elongation, tapering, and curvature. This observation suggests that these dimensions are either strongly correlated with or actually correspond to some of those used by the visual recognition system for the encoding and/or representation of visual shapes. Indeed, had this not been the case, the representation of the structural relationships between items making a conjunction set would have been misaligned relative to the main axes of the *psychological* shape space used by the visual recognition system. Two possible forms of such a misalignment are

illustrated in Fig. 9. In these examples, it can be seen that each item of a particular set has a unique value (i.e. projection) on at least one of the dimensions of *psychological* shape space, therefore reducing the conjunction task to a single-dimension one. That is, with the kinds of psychological representations of structural relationships illustrated in Fig. 9, consideration of stimulus properties along only one shape dimension is sufficient to specify any particular item uniquely, therefore bypassing the need for shape integration. Given the marked performance contrast observed in ELM between single-dimension and conjunction sets as well as the type of errors he made with the latter stimulus sets, it is quite clear that the conjunction task used here could not be reduced to one involving a single dimension. This strongly argues for the psychological validity of the dimensions of elongation, tapering, and curvature used by our shape generation

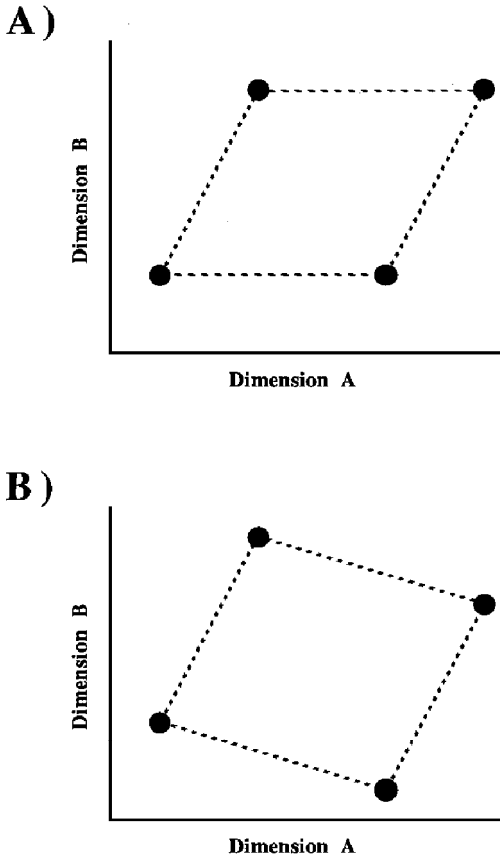


FIG. 9. Hypothetical structural relationships, in psychological shape space, between conjunction items if variations in stimulus properties had been misaligned relative to the psychological dimensions by which shapes are represented.

algorithm. As noted previously, these dimensions are closely related to some that have previously been suggested by theoretical approaches to shape perception and representation (Barr, 1981; Biederman, 1987; Brooks, 1981; Pentland, 1985, 1988). The observations of Experiment 4 provide empirical support for these suggestions.

## EXPERIMENT 5

The previous experiment has provided conclusive evidence for a shape integration impairment in ELM. It remains to be determined, however, whether this impairment affects operations involved in gaining an internal description of the input or whether it concerns the retrieval of stored structural knowledge about objects. The preliminary testing conducted on ELM, which was reported above, has indicated that the patient is impaired at the retrieval of stored structural knowledge about biological objects, whereas no deficit was noted with respect to the encoding of properties of visual objects. Although suggestive, these observations were not specifically tied to the question of shape integration, and the exact relationship between ELM's shape integration impairment and his category-specific visual agnosia remains to be established. An additional experiment is therefore required to determine the processing stage at which shape integration fails in ELM. Experiment 5 focused on integration for shape perception and used a task where a target had to be matched, by pointing, to another identical object in a set of four stimuli presented simultaneously.

In several respects, Experiments 4 and 5 were very similar. Thus, the same stimulus sets were used, and targets had to be discriminated from the other objects of the set to which they belonged. In addition, the task used in Experiment 5 was again entirely non-verbal, and the response required of the subject was one of pointing. The critical difference between Experiments 4 and 5 concerned the requirements in terms of storage of shape properties. Experiment 4 conformed to the usual requirements of a recognition task, where a set of different objects and their associated experiences and responses had to be retained in memory. On any trial, the target displayed had to be referred to a unique object of the set held in memory and then mapped to the appropriate assigned location. In contrast, Experiment 5 essentially involved perceptual matching and therefore implied only a minimal contribution of memory for shape. First, the target shape had to be held in memory over a very brief interval (1 sec). Second, and more importantly, the storage of the target did not involve a contrast between its properties and those of other objects of the same set held in memory. Rather, the unique shape to be held in memory on any given trial had to be contrasted to actual stimuli that were visible to the subject.

From this, it appears that if ELM's integration impairment is manifest in the encoding of visual shapes, the patient should show the same kind of performance pattern in the present experiment as he did in Experiment 4. In contrast, if shape

integration fails at the level where stored knowledge of shapes is retrieved for perceptual categorization, Experiment 5 should not reveal the dimensionality effect observed in ELM in the previous experiment.

## Methods

The stimulus sets tested were the same as those used in Experiment 4. The complete experiment comprised 8 blocks of 60 trials each. Every item of every shape set was tested a total of 20 times. The order of trials within each block was random. On each trial, one target shape was first presented at the centre of the display screen. It was then removed, and after a delay of 1 sec the four stimuli of the set to which the target belonged were presented simultaneously, each at the one corner of the screen (random shape/location assignments). The subject's task was to point to the target shape that had just been shown, and he could take his time to respond.

## Results and Discussion

ELM made only one (0.004%) error with targets taken from conjunction sets. In contrast, with targets taken from single-dimension sets, he committed an error on 8.8% of trials, which is significantly greater than the error rate found with conjunction sets,  $\chi^2(1) = 19.1$ ,  $p < .001$ . This dimensionality effect is congruent with the greater perceptual similarity between items of single-dimension sets and provides no evidence for a deficit in shape integration with respect to perceptual encoding in ELM.

The performance of ELM in the visual matching task of the present experiment is in direct opposition to his results in the visual recognition paradigm of Experiment 4, even though the perceptual discriminations that had to be performed were identical. This crossover interaction between ELM's results in Experiments 4 and 5 clearly indicates that the shape integration impairment he exhibited in a recognition task cannot be due to a failure of perceptual encoding. Rather, the nature of the difference between the demands of each experiment indicates that this shape integration impairment must be attributed to a deficit in retrieving structural knowledge about objects. Specifically, it is concluded that the patient's visual recognition deficit is characterized by the partial retrieval of shape information from memory.

It has been underlined in the discussion of Experiment 4 that the failure of shape integration in ELM has implications about the occurrence of such a process in normal object recognition. Given the additional information provided by the results in Experiment 5, it is now possible to restate these implications in a more precise way. First, the fact that brain damage can cause a deficit such that only a subset of an object's shape features can be retrieved from structural knowledge means that this knowledge is stored in a distributed fashion, with discrete features representing the shape properties of the object along a number

of separate dimensions. Second, and perhaps more importantly, the present results indicate that visual recognition must be conceived as an active process involving an assembly operation through which the discretely stored features representing the structural knowledge of an object are integrated.

## EXPERIMENT 6

The study of shape integration in ELM was initially motivated by the properties of his confusion errors when attempting to recognize line drawings of fruits and vegetables (Experiment 1). Ultimately, the goal of the present study is to propose an explanatory hypothesis for the patient's agnosic symptoms—in particular, its category specificity. The previous experiments have been successful in establishing the existence of a shape integration impairment in ELM and in providing evidence as to the level of processing at which it is manifest. The major part of this evidence, however, has come from experiments in which the stimuli used had no semantic content and from which even verbal processing was excluded (cf. Experiments 4 and 5). In these respects, our study of shape integration in ELM has become quite remote from the actual manifestations of his agnosia. It therefore remains to be determined whether the patient's deficit in structural knowledge integration relates to the selectivity of his agnosia for biological objects. Experiment 6 attempts to provide this missing link.

To be able to show that an impairment of structural knowledge integration is an essential element for an account of the category-specificity of ELM's visual agnosia, two demonstrations need to be made: (1) In attempting to recognize biological objects, the patient exhibits unequivocal signs of a shape integration deficit. (2) In attempting to recognize man-made objects, ELM *does not* show signs of a shape integration deficit.


With respect to the first required demonstration, signs of a shape integration impairment for real biological objects were observed in Experiments 1 and 2, but some degree of ambiguity remained in both sets of data. The hypothesis of a shape integration impairment emerged as a post hoc attempt to explain the picture–word matching data obtained in Experiment 1. In the second experiment, the error pattern observed was congruent with a shape integration impairment. However, the pictures used on the negative trials of this picture–word matching task most often did not correspond to any known object. From this, it could be argued that the testing situation has little environmental validity in that, on most trials, there was no possibility for confusions to occur between pairs of items that could be mapped to stored structural *and* semantic knowledge. Additional demonstration for a shape integration deficit in the recognition of biological objects is therefore needed.

The question becomes markedly more complex as far as the second required demonstration is concerned. Indeed, it is quite obvious that man-made objects typically share some of their constituent shape features with other objects, and

therefore that the processing of several of their visual properties is required for unique identification. Why then should we expect a shape integration impairment to occur with biological objects but not with artefacts? In other words, what is the difference between artefacts and biological objects that would allow ELM to perform shape integration successfully with the former but not the latter? At this point, it seems that an appeal to an additional explanatory factor must be made. Specifically, it appears that, if shape integration is to be involved in an account of the category specificity of ELM's agnosia, it must be assumed that this operation is modulated by another factor along which biological objects and artefacts differ. Relevant information on this question may be found in the work of Humphreys et al. (1988).

As reported previously, these authors have studied the structural relationships between semantically related objects of various categories by gathering subjective similarity ratings produced by normal individuals. What their results have shown is that, as a general rule, semantically related artefacts tend to be visually dissimilar from one another, whereas the opposite is true for biological objects. An extrapolation of this result would suggest that visually similar (and possibly structurally related in terms of shape features) pairs of man-made objects would usually be more semantically distant (i.e. belong to separate semantic categories) from each other than structurally related pairs of biological objects. If it is further assumed that semantic distance between visually related objects may somehow (possibly by some form of feedback) contribute to the integration of structural knowledge about them, and a brain-damaged patient whose deficit

### SINGLE DIMENSION



### CONJUNCTION




FIG. 10. Stimuli used for the single-dimension and conjunction sets in Experiment 6.

concerns solely this operation might show intact visual recognition of artefacts, even though the visual recognition of biological objects is severely impaired. Experiment 6 will test this hypothesis.

One "single dimension" set and one "conjunction" set were used in this experiment (Fig. 10), and these were generated according to the same principles as those described in the method section of Experiment 4. Each of the shapes could alternatively be referred to as fruits or vegetables or as artefacts. Their recognition was tested in ELM using a naming paradigm.

Initially, it was expected that ELM would make numerous errors in this task, especially with items from the conjunction set. Much to our surprise, however, it was quickly evident that the patient made very few if any errors. This was true even with the stimuli from the conjunction set referred to as fruits or vegetables. This observation is in marked contrast with the previous results indicating that ELM is severely impaired: (1) in the recognition of fruits and vegetables, and (2) with conjunction sets of shapes with no semantic content. It seems that two factors may have contributed to the patient's improved performance in the task. First, contrary to the previous experiments, which showed impaired recognition of fruits and vegetables in ELM (Experiments 1 and 2), both the set of possible shapes that could be presented as well as the number of object categories used were very restricted. In these respects, the present experiment has features in common with Experiment 3, where ELM showed a good level of performance even with items that had been frequently confused when presented among a larger set of objects (cf. Experiment 1). These points converge in showing a strong dependency of ELM's visual recognition performance on the extent of the contrastive set from which a given visual target needs to be discriminated. Second, and more importantly here, the essential methodological contrast between Experiments 4 and 6 is that the stimuli used in the latter were given a semantic content by referring to them as real-world objects. It would seem, then, that the contrast between ELM's performance in Experiments 4 and 6 provides preliminary evidence for a contribution of the semantic properties of objects in supporting the integration of shape knowledge.

As error rate could not serve as a dependent variable in Experiment 6, the task used was one of speeded naming, where the main measure of performance was response time (RT). From the results of the previous experiments, which provided evidence for an impairment of the shape knowledge integration in ELM, it was predicted that his RTs would be higher with items taken from a conjunction set than with items from a single-dimension set. However, from the hypothesis of a semantic modulation of the process of integration of shape knowledge, it was predicted that the dimensionality effect would be present when items are referred to as fruits or vegetables, but not when they are referred to as artefacts. Indeed, as noted above, fruits and vegetables are obviously related to one another semantically by virtue of the fact that they belong to the same semantic category. In contrast, the artefacts selected for Experiment 6 are not, or only remotely, semantically related.

## Methods

Two stimulus sets were generated (Fig. 10). One was made of shapes differing from one another on elongation (single-dimension set), with shape values on curvature and tapering remaining constant across items. The other was a conjunction set, and shape properties varied on the dimensions of elongation and tapering, with stimulus values remaining constant on the dimension of curvature. Each of the stimuli used could serve as a recognizable depiction of a real fruit or vegetable (single dimension: asparagus, cucumber, melon, and orange; conjunction: cucumber, carrot, melon, and pear) or of a man-made object (single dimension: pen, cigar, balloon, and baseball; conjunction: cigar, tentpeg, balloon, and spinning top). Each stimulus set was tested twice in blocks of 52 randomly ordered trials each, once referring to items as fruits or vegetables, once referring to them as artefacts. Before each block, the subject was told that he would be shown, on any trial, one of four shapes, and he was given the names of the stimuli tested. He was instructed to name the target as rapidly as possible while avoiding errors. Trials were preceded by a 500-msec fixation point. The target was then shown by itself at the centre of the display screen, and responses were registered by a voice-key interfaced with the control computer. The experimenter entered the subject's response after each trial, and the subject was given feedback about the accuracy of his responses.

## Results and Discussion

ELM's correct RTs in each condition of Experiment 6 are illustrated in Fig. 11, and error rates are presented in Table 4. The correlation between correct RTs and error rates was of + 0.93, therefore indicating no speed-accuracy trade-off.

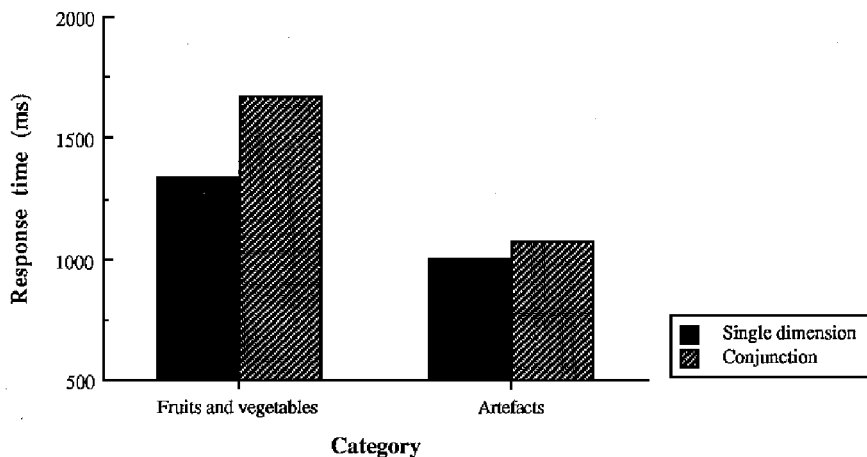


FIG. 11. Average response times in each condition of Experiment 6.



TABLE 4  
 Error Rates in the Speeded Naming of Shapes Taken  
 from Single-dimension or Conjunction Sets, Referred  
 to as Fruits or Vegetables or as Artefacts

<i>Reference Category</i>	<i>Stimulus Set</i>	<i>%</i>
fruits/vegetables	single dimension	3.9
fruits/vegetables	conjunction	9.6
artefacts	single dimension	1.9
artefacts	conjunction	3.9

Correct RTs were analysed with a two-way ANOVA, with dimensionality of the stimulus set (single dimension vs. conjunction) and semantic category (fruits/vegetables vs. artefacts) treated as between factors (i.e. each RT is considered as an independent sample of the subject's performance in a particular condition). This analysis revealed a main effect of semantic category  $F(1, 164) = 78.6$ ,  $p < .001$ , with longer RTs for fruits and vegetables than for artefacts. The main effect on dimensionality was also significant,  $F(1, 164) = 14.6$ ,  $p < .005$  and indicated longer RTs with the conjunction set than with the single-dimension set. These main effects were qualified by a Dimensionality  $\times$  Category interaction,  $F(1, 164) = 6.2$ ,  $p < .02$ . Simple effects of this interaction indicated that the effect of dimensionality was significant for fruits and vegetables,  $F(1, 164) = 18.6$ ,  $p < .001$ , but not for artefacts,  $F(1, 164) = 1$ . ELM's error rates were rather low (Table 3), and there was no significant effect of condition on accuracy,  $\chi^2(3) = 3.8$ , n.s.

The results of Experiment 6 provide clear evidence that ELM's structural knowledge integration impairment, which was documented in the previous experiments, is modulated by the semantic properties of the objects presented. Thus, with the fruit and vegetable labels, the patient took longer to recognize items taken from a conjunction set than items from a single-dimension set. This is explained by the fact that whereas processing of information about a single shape dimension was sufficient for the unique identification of targets from the single-dimension set, integration of features along two separate shape dimensions was required with items from the conjunction set. The effect of dimensionality of the shape sets did not occur, however, when items were referred to as artefacts. This contrast across semantic category conditions on the dimensionality effect can only be explained by the semantic properties attributed to the items, because exactly the same shape sets were used in the fruit/vegetable and artefact conditions. Two kinds of semantic properties may be held responsible for the modulation of the effect of dimensionality of the shape set in Experiment 6.

One is simply that items were referred to as man-made objects in one semantic condition and as biological objects in another. It could be argued, for instance, that these two broad categories differ somehow from one another on the way semantic representations are organized. In turn, this difference in terms of semantic organization could determine some aspect of the structural knowledge integration process such that, after brain damage, this operation might be impaired for one category (biological objects) but not for another (artefacts). Although this hypothesis cannot be rejected by the present observations, it comprises some problematic elements. Thus, this explanation is essentially a restatement of the empirical dissociation observed in CSVA for biological objects. It says nothing about the organization of the knowledge systems that contribute to visual object recognition apart from the introduction of special functional rules that are difficult to justify beyond the fact that they thoroughly fit the clinical manifestations of the disorder.

Another account of the semantic modulation of the structural knowledge integration process that appears more plausible rests on the fact that whereas the fruit and vegetable labels used referred to semantically related objects, artefact labels referred to objects that are semantically distant from one another. Although the exact mechanisms remain to be determined, it appears possible that the integration of structural knowledge might be facilitated by semantic distance between objects. The past work of Humphreys et al. (1988), which showed that the visual recognition of semantically related and visually similar objects is particularly sensitive to the effects of brain damage, is congruent with this hypothesis.

An empirical disambiguation of the two alternative hypotheses described here to account for the semantic modulation of structural knowledge integration in ELM would have required the study of his performance with a conjunction set of semantically related artefacts. Indeed, in this case, it would have been possible to dissociate the effect of semantic proximity from that of referring to artefacts. Unfortunately, we were unable to generate this kind of stimulus set with the shape dimensions of elongation, tapering, and curvature. It remains, however, that, as a general rule CSVA for biological objects does not entirely spare the recognition of artefacts and the recognition of particular man-made categories is just as affected as is the recognition of biological objects (Damasio, 1990; Farah, 1991; see also footnote 1) This is also true for ELM. Thus, the patient mentioned difficulties in recognizing makes of car, and he has to put a flag on his own car in order to distinguish it from others in parking lots. The patient also exhibits difficulties in the visual recognition of musical instruments. He has expert knowledge of brass instruments and is eventually able to name them correctly when they are presented as line drawings. However, this is possible only after a lengthy examination of the item, and the patient spontaneously indicates that these correct identifications are based on local properties of the objects (e.g. number of keys, mouthpiece, etc.). For musical instruments

other than brass, ELM is as severely agnostic as he is for biological categories, and he may sometimes make gross confusions (e.g. a banjo called “a guitar”). These observations suggest that it is not sufficient for an item to belong to artefact categories to prevent its impaired visual recognition in CSVA for biological objects.

In closing, one possible limitation to our interpretation of the results in Experiment 6 needs to be discussed. Thus, although the conjunction condition used involved processing of shape features on the dimensions of elongation and tapering, only elongation was tested in the single-dimension condition. Attempts were made at finding sets of fruits and vegetables and of man-made objects that represent shape variations on tapering only to constitute an additional “single-dimension” control condition. Unfortunately, such sets could not be found. The reason why the lack of a “tapering-only” condition might be construed as problematic is that one may assume that the semantic effect observed with conjunction sets is unrelated to the feature integration required by this condition but is, rather, linked to the necessity of processing tapering information. Under this assumption, then, it would be predicted that performance in a tapering-only condition would have been as strongly influenced by semantic category as was observed with conjunction sets. Although the present paper does not report any empirical fact that can directly disprove this alternative account, we should note that it assumes a special status for the shape dimension of tapering in the way it interacts with semantic knowledge. Apart from the fact that the principle form which such a special status would originate is far from clear, Experiment 3 suggests that tapering does not constitute a special problem for ELM when a single-dimension task involving semantic content is presented (see Fig. 6). Perhaps more crucially, a separate series of experiments conducted in collaboration with Mike Dixon, which are the object of a separate report, clearly falsifies the notion that the semantic effect observed here in the conjunction condition is related to the use of tapering as a critical dimension. These experiments have used a name-learning paradigm for single-dimension or conjunction sets of shapes similar to those that served here. Contrary to Experiment 6, however, the choice of shapes to be learned was unconstrained by their resemblance to real-world objects. What the results of these experiments show is that although semantic proximity has a dramatic effect on ELM’s *accuracy* with conjunction sets, it has no effect on single-dimension sets even when they involve tapering.

## GENERAL DISCUSSION

The present experimental series has provided evidence for two main points with respect to visual object recognition: (1) structural knowledge about objects is stored as sets of discrete features that need to be integrated for unique visual identification, and this integration process may be impaired following inferior

temporal lobe damage; (2) a deficit of structural knowledge integration may be a critical component in the explanation of CSVA for biological objects. These aspects of the current investigation will be discussed separately, and we wish to emphasize that these two issues should not be confused, as they are partially independent from one another. In particular, one point we have shown here is the existence of a strong correlation between the failure of a brain-damaged patient to perform structural knowledge integration and CSVA. It is argued below that this correlation reflects a causal relationship between impaired structural knowledge integration and CSVA. Whether or not this proposed causality is supported by future research should have no bearing on the demonstration of a shape integration failure in ELM and on its implications for theories of normal visual object recognition.

## Structural Knowledge Integration

Several pieces of evidence obtained from ELM indicate a deficit of shape integration. Thus, in matching graphic depictions of fruits and vegetables to spoken names, the patient systematically confused objects that seemed to share critical shape properties (Experiment 1), and his error rate on negative trials increased systematically with the number of shared features between the target object and the picture displayed (Experiment 2). Whereas the patient could adequately perform object classifications based only on discriminations of elongation, curvature, or tapering (Experiment 3), his visual recognition performance was markedly affected by an increase in the dimensionality of the stimulus set when it led to a requirement for integration of shape properties for unique identification (Experiment 4). This ensemble of observations from ELM, along with other less critical aspects of the data that were reported earlier, points to a failure in processing multiple shape dimensions for visual object recognition; that is, a shape integration impairment.

Experiment 5 has been critical for a specification of the functional locus of this impairment and has shown that it specifically affects the retrieval of structural knowledge about objects. Thus, a perceptual matching task with minimal requirements in terms of memory for shape has shown a reversal of the effect of the dimensionality of the stimulus set, with significantly better performance with conjunction than with single-dimension sets. This led to the conclusion that ELM's shape integration deficit specifically affects the retrieval of structural knowledge and not the encoding of the properties of the visual stimulus.

The observation of a deficit such as that reported in ELM as well as the elaboration of the concept of structural knowledge integration are, to our knowledge, unprecedented in the literature on visual object recognition. However, in different contexts, two most notable uses of the notion of integration have been made by other groups of investigators to refer to some aspect of visual processing. It is important here to specify the way in which these proposals differ from ours.

Treisman and her collaborators (as well as numerous others who have followed up on her work—e.g. Treisman & Gelade, 1990; Treisman & Sato, 1990) have studied feature integration with respect to visual encoding in normal observers. What Treisman has proposed is that visual domains such as colour and line orientation are processed via separate pathways and that the adequate *perception* of conjunctions of separable features in a particular object requires integration through spatial attention. In a separate series of investigations, this time in a brain-damaged patient with bilateral occipital damage (HJA), Humphreys and Riddoch (1987a; Riddoch & Humphreys, 1987b) have reported a deficit in integrating local elements of visual objects into a coherent perceptual whole and in linking this local information with global aspects of shape. In contrast to his subtly impaired processing of visual input, however, HJA retained good stored knowledge of the visual properties of the objects. Moreover, and strikingly, HJA has preserved his capacity to encode global information about the shapes of objects—that is, his deficit is specific to the integration of local/separable elements of visual stimuli. Such observations have led Humphreys and Riddoch to propose the term “integrative visual agnosia” to describe their patient’s impairment.

Certainly the most crucial distinction that may be drawn between the concept of structural knowledge integration proposed here and the theories put forward by Treisman and collaborators and by Humphreys and Riddoch is the level of processing that is addressed. Whereas these authors have discussed integration with respect to the encoding of visual images, our proposal concerns the process by which knowledge stored in memory about the visual properties of objects is retrieved. Another important contrast between the different theories relates to the properties that are subjected to the integration process. The main focus of Treisman’s work concerned the integration of visual features that are encoded by separate processing modules. In the case of HJA, the problem was related to the integration of spatially disparate (although often connected) visual elements. HJA had apparently no impairment with the representation of global form. In contrast, what we have studied in ELM is the integration of global shape features—that is, the features studied are not spatially isolable from the global form. Whether or not ELM’s deficit may extend to the integration of visual knowledge about separable visual domains (e.g. conjunctions of colour and line orientation) or about separable elements of a visual object (e.g. the way in which pairs of lines are joined) has yet to be investigated.

Beyond the patient studied here, there are grounds for believing that the current findings may have important implications for the operation of the normal visual recognition system. We suggest that these implications, described later on, will need to be addressed in the future by cognitive theories of visual recognition and that they may affect concepts about the functional properties of the inferior temporal lobes.

As stated previously, the fact that brain damage may lead to a failure of structural knowledge integration must imply that this process contributes to visual object recognition in normal individuals. So far, cognitive models of visual recognition have mainly focused on the content of structural knowledge—that is, the dimensions on which shape is represented—rather than on retrieval procedures. In that respect, it does not appear that the discovery of a process of structural knowledge integration is directly incongruent with any of these theories. A common and most often implicit assumption in such models, however, seems to be that visual object representations are stored as integral units—that is, virtual objects—that only need to be activated by the appropriate input for perceptual categorization to occur. This assumption is not supported by the present observations, which suggest the need of an integration process for “access to” or “retrieval of” structural knowledge and, by extension, that visual object representations are stored as sets of discrete features rather than as integral units. It appears that these facts will need to be incorporated into future models of visual object recognition. This theoretical shift may, in fact, have practical consequences for the study of visual recognition. Indeed, provided some experimental paradigm that is sensitive to the occurrence of structural knowledge integration in normal individuals can be developed, it should be possible to conduct studies of the dimensions under which object knowledge is represented—that is, the content of those representations. For instance, this type of study could take a form comparable to that of Experiment 4, where the dramatic effect of dimensionality of the stimulus set in ELM led us to propose that elongation, curvature, and tapering are dimensions under which visual shapes are represented.

In addition to purely cognitive aspects of visual recognition, it is interesting to note that ELM's impairment following bilateral inferior temporal lobe damage is congruent with current knowledge, mainly gathered from the monkey, about the function of this area (for a review of behavioural, anatomical, and physiological findings, see Grüsser & Landis, 1991, Chapter 7). In particular, a wealth of studies have shown that monkeys with bilateral inferior temporal lobectomies suffer from a profound visual recognition impairment that seems specifically linked to a memory deficit for the visual properties of objects (for a review, see Mishkin, 1982). These findings largely parallel the observations made in ELM. Furthermore although it is not yet entirely clear how the concept of structural knowledge integration can be applied to the visual function of the inferior temporal cortex, an electrophysiological study by Gochin, Miller, Gross, and Gerstein (1991) has provided observations consistent with this notion. These authors have studied the visual responses of individual cells in the inferior temporal cortex (area TE) using multi-electrode recording and applying auto- and cross-correlational analyses to the obtained spike trains. The findings of greater importance, in the present context, which were reported by Gochin et al. (1991) are the following: (1) Nearby neurones show greater similarity in their stimulus selectivity than do neurones that are distant from one another,

suggesting the existence of patches of neurones of a few hundred  $\mu\text{m}$  in diameter that share common response properties. (2) The rate of direct connections between inferior temporal units declines only slightly as inter-neurone distance increases, thus suggesting that intrinsic connections within the inferior temporal cortex distribute information from the relatively localized neurone patches over to more distant regions. (3) Inferior temporal neurones were shown to be driven by 39 inputs, on the average; this is considerably larger than the range of about 2 to 10 inputs for neurones in the striate cortex (Creutzfeldt & Ito, 1968; Michalski, Gerstein, Czarkowska, & Tarnecki, 1983; Toyama, Kimura, & Tanaka, 1981). Overall, these observations suggest that the inferior temporal cortex is an area endowed with a rich connectivity where long-range interactions between units responding to different stimulus features may occur. These properties are consistent with what we may expect of a brain region involved in the integration of the multiple and separately represented properties of a particular object.

## Visual Agnosia to Biological Objects

Several pieces of circumstantial evidence suggest that ELM's structural knowledge integration impairment may be a causal factor for his CSVA.

First, the hypothesis of a deficit in shape integration in ELM was initially derived from the error pattern he exhibited in a picture-word matching task where the stimuli presented were line drawings of common fruits and vegetables (Experiment 1). In other words, one of the most prominent features of the agnostic manifestations seen in the patient is entirely congruent with the possibility that his demonstrated impairment for shape integration plays a critical role in his failure to visually recognize biological objects.

Second, the background information gathered from the patient, which was presented in the first section of this paper (see also Decter et al., submitted), indicates that the functional locus of the impairment responsible for his agnosia concerns the representation or the retrieval of structural knowledge about objects. Similarly, Experiment 5 has shown that ELM's deficit for shape integration specifically affects the retrieval of stored structural knowledge.

Third, Experiments 2 and 4 have shown that individual shape dimensions seem to be processed by ELM with different degrees of priority in visual recognition tasks that require shape integration. In particular, it was noted that confusions between items differing from one another on elongation were markedly less frequent than confusions between items having a common degree of elongation but differing from one another on either curvature or tapering. It seems possible that this variation across shape dimensions may be due to differences in the degree of attention allocated to individual dimensions (Nosofsky, 1986), which, in turn, may relate to differences in the perceptual salience of particular dimensions for the processing of structural contrasts between the items

presented. More importantly, the evidence from Experiment 1 indicates quite clearly that a comparable order prioritization occurred with pictures of fruits and vegetables. Specifically, it may be noted that items differing on elongation were very seldom confused and that errors mainly involved other aspects of shape. This may be seen more easily from Fig. 3, where one group of elongated objects emerged very distinctly from other, more rounded objects. These elongated items were often confused with one another despite marked difference between them on curvature or tapering; they were practically never confused with the rounded objects of the set used.

However, the most direct link between ELM's structural knowledge integration deficit and his CSVA was provided by Experiment 6, where the patient had to name stimuli taken from single-dimension or conjunction sets that were referred to as either fruits or vegetables or as artefacts. Indeed, as discussed earlier, if ELM's structural knowledge integration deficit was to apply in the explanation of the category specificity of his agnosia, it needed to be shown that the deficit is manifest for biological objects but not for artefacts. In other words, it was necessary to demonstrate that semantic properties of the items modulate the occurrence of symptoms of a shape integration impairment in the patient. This was achieved here by showing that ELM exhibits a large increase in naming times with an increase in the dimensionality of the stimulus set when items were referred to as fruits or vegetables. In contrast, with precisely same shape sets, dimensionality had no effect if items were referred to as artefacts. From these observations, it was proposed that semantic distance between structurally related items somehow contributes to the structural knowledge integration process such that an isolated shape integration deficit will only be manifest with items that do not possess semantic features that are radically different. This explanation of CSVA for biological objects in terms of structural knowledge integration deficit modulated by semantic properties of the objects is entirely congruent with (and in fact the hypothesis was initially derived from) the observation of Humphreys et al. (1988) that semantically related biological objects are more structurally similar to one another than are semantically related artefacts.

To summarize, what the above arguments mean is that any patient who suffers from a structural knowledge integration impairment but whose semantic memory is relatively spared will show CSVA for biological objects. Of course, there are still a number of points that remain to be verified to prove this hypothesis ultimately—the most important among them are a stronger demonstration that semantic relatedness affects the operation of structural knowledge integration (evidence for which was briefly presented in the discussion of Experiment 6 and will be the object of a separate report) and observations pertaining to the uniqueness of the assumed link from impaired structural knowledge integration to CSVA for biological objects. Another point that must be emphasized here is that the account proposed for ELM's agnosia should not be taken as a general explanation of CSVA for biological objects. Indeed, it is quite clear that patients



such as JB and KR (presented in the Introduction) do not have impaired structural knowledge integration, given that they can perform reality decisions accurately. It should also be noted that category-specific deficits in general can originate from a variety of impairments, as discussed in the Introduction, which have nothing to do with the retrieval of structural knowledge.

The current findings as well as the conclusions derived from them relate in interesting ways to some of the previous accounts of CSVA for biological objects. Along with Etcoff et al. (1991), Sartori and Job (1988), and Sartori et al. (1992), we propose that the disorder may be the result of a failure in retrieving structural knowledge from memory. However, although these previous investigators argued for a loss of structural knowledge in CSVA, the present results suggest that structural knowledge is preserved and that the critical deficit specifically concerns its access or retrieval. The fact that ELM's confusion errors for fruits and vegetables are clearly organized around a principle of shared structural features between objects (Experiments 1 and 2) indicates that the patient maintains some structural knowledge about the items he fails to recognize. His adequate classification performance (Experiment 3) when discriminations between fruits and vegetables involved a single-dimension contrast (either elongation, curvature, or tapering) also argues for preserved structural knowledge about biological objects. What was shown here is, rather, that ELM can effect only a partial retrieval of his retained structural knowledge, such that this retrieval fails to specify uniquely one particular object against its structural neighbours—that is, other items with common features.

It should also be pointed out that the results reported here suggest that the hypothesis of Sartori and Job (1988), Sartori et al. (1992), and Etcoff et al. (1991) about a division of structural knowledge as a function of semantic category is not necessary to account for CSVA for biological objects. Thus, the results of Experiment 4 show that the deficit responsible for ELM's failure to retrieve structural knowledge from memory is quite general and applies even to items to which no semantic content was attributed. Rather, the results of Experiment 6 show that a semantic modulation of the structural knowledge integration process is a more probable explanation for category specificity and for the associated deficits affecting the recognition of specific classes of man-made objects.

Of the various previous explanations of CSVA for biological objects, the one proposed here seems closest to that of Riddoch and Humphreys (1987a; Humphreys & Riddoch, 1987b; Humphreys et al., 1988) in its emphasis on the facts that: (1) it is possible to account for the disorder without assuming category-specific operations, and (2) the nature of structural relationships between semantically related biological and man-made objects is a critical element for the explanation of the disorder. The fundamental advance provided by the present experimental series with respect to the past work of these authors on CSVA is twofold. First, it has allowed the identification of the operation—structural knowledge integration—responsible for the structural similarity constraint

that Riddoch and Humphreys have shown to act on CSVA. Second, it has shown that this structural similarity constraint does not actually concern overall visual proximity but, rather, the existence of shared shape features between objects.

## CONCLUSIONS

The study of a patient with CSVA for biological objects following bilateral inferior temporal lobe lesions has allowed the identification of a deficit of structural knowledge integration. It is argued that the discovery of this kind of impairment in a brain-damaged patient has fundamental implications for cognitive theories of visual object recognition and that it may contribute to the specification of the visual function of the inferior temporal cortex. The finding, in Experiment 6, of a modulation of the structural knowledge integration process in the patient studied by the semantic properties of the objects has provided evidence for the applicability of the current set of observations to an explanation of the phenomenon of CSVA for biological objects.

## REFERENCES

- Barr, A.H. (1981). Superquadrics and angle-preserving transformations. *IEEE Computer Graphics and Applications*, *1*, 1–20.
- 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.
- Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, *94*, 115–147.
- Brooks, R.A. (1981). Symbolic reasoning among 3-D models and 2-D images. *Artificial Intelligence Journal*, *17*, 285–348.
- Caramazza, A., & McCloskey, M. (1988). The case for single-patient studies. *Cognitive Neuropsychology*, *5*, 517–528.
- Creutzfeldt, O., & Ito, M. (1968). Functional synaptic organization of primary visual cortex neurons in the cat. *Experimental Brain Research*, *6*, 324–352.
- Damasio, A.R. (1990). Category-related recognition defects as a clue to the neural substrates of knowledge. *Trends in Neuroscience*, *13*, 95–98.
- Decter, M., Bub, D., & Chertkow, H. (submitted). Multiple representation of object concepts: Evidence from category-specific agnosia. *Cognitive Neuropsychology*.
- 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. (1991). Patterns of co-occurrence among associative agnosias: Implications for visual object representation. *Cognitive Neuropsychology*, *8*, 1–19.
- 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.
- Gaffan, D., & Heywood, C.A. (1993). A spurious category-specific visual agnosia for living things in normal human and nonhuman primates. *Journal of Cognitive Neuroscience*, *5*, 118–128.
- Gochin, P.M., Miller, E.K., Gross, C.G., & Gerstein, G.L. (1991). Functional interactions among neurons in inferior temporal cortex of the awake macaque. *Experimental Brain Research*, *84*, 505–516.
- Grüsser, O.-J., & Landis, T. (1991). *Visual agnosias and other disturbances of visual perception and cognition*. Boston, MA: CRC Press.
- Hart, J., Berndt, R.S., & Caramazza, A. (1985). Category-specific naming deficit following cerebral infarction. *Nature*, *316*, 439–440.
- Hart, J., & Gordon, B. (1992). Neural subsystems for object knowledge. *Nature*, *359*, 60–64.
- Hillis, A.E., & Caramazza, A. (1991a). Category-specific naming and comprehension impairment: A double-dissociation. *Brain*, *114*, 2081–2094.
- Hillis, A.E., & Caramazza, A. (1991b). Mechanisms for assessing lexical representations for output: Evidence from a category-specific semantic deficit. *Brain and Language*, *40*, 106–144.
- Hoffman, D.D., & Richards, W. (1984). Parts of recognition. *Cognition*, *18*, 65–96.
- Humphreys, G.W., & Riddoch, M.J. (1987a). *To see but not to see: A case study of visual agnosia*. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.
- Humphreys, G.W., & Riddoch, M.J. (1987b). 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*, *85*, 67–103.
- Lockhead, G.R. (1970). Identification and the form of multi-dimensional discrimination space. *Journal of Experimental Psychology*, *8*, 1–10.
- Loomis, J.M. (1982). Analysis of tactile and visual confusion matrices. *Perception and Psychophysics*, *31*, 41–52.
- Marr, D. (1982). *Vision*. San Francisco, CA: Freeman.
- 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.
- McCarthy, R.A., & Warrington, E.K. (1988). Evidence for modality-specific meaning systems in the brain. *Nature*, *334*, 428–430.
- 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.
- Michalski, A., Gerstein, G.L., Czarkowska, J., & Tarnecki, R. (1983). Interactions between cat striate cortex neurons. *Experimental Brain Research*, *51*, 91–107.
- Mishkin, M. (1982). A memory system in the monkey. *Philosophical Transactions of the Royal Society of London*, *B298*, 85–95.

- 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.
- Palmer, S.E. (1977). Hierarchical structure in perceptual representation. *Cognitive Psychology, 9*, 441-474.
- Pentland, A. (1985). On describing complex surface shapes. *Image and Vision Computing, 3*, 153-162.
- Pentland, A. (1988). *Automatic extraction of deformable part models*. Vision Sciences Technical Report no. 104, MIT Media Laboratory.
- Pietrini, V., Nertempi, P., Vaglia, A., Revello, M.G., Pinna, V., & Ferro-Milone, F. (1988). Recovery from herpes simplex encephalitis: Selective impairment of specific semantic categories with neuroradiological correlation. *Journal of Neurology, Neurosurgery, and Psychiatry, 51*, 1284-1293.
- Pinker, S. (1984). Visual cognition: An introduction. *Cognition, 18*, 1-63.
- 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.
- Riddoch, M.J., & Humphreys, G.W. (1987a). Visual object processing in optic aphasia: A case of semantic access agnosia. *Cognitive Neuropsychology, 4*, 131-185.
- Riddoch, M.J., & Humphreys, G.W. (1987b). A case of integrative visual agnosia. *Brain, 110*, 1431-1462.
- Sacchet, 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, 5*, 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., Guistolisi, 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.
- Sirigu, A., Duhamel, J.-R., & Poncet, M. (1991). The role of a sensorimotor experience in object recognition. *Brain, 114*, 2555-2573.
- Snodgrass, J.G., & Vanderwart, M.A. (1980). Standardized set of 260 pictures: Norms for name agreement, image agreement, familiarity, and visual complexity. *Journal of Experimental Psychology: Human Learning and Memory, 6*, 174-215.
- Toyama, K., Kimura, M., & Tanaka, K. (1981). Cross-correlation analysis of interneuronal connectivity in cat visual cortex. *Journal of Neurophysiology, 46*, 191-201.

- Treisman, A., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology, 12*, 97–136.
- Treisman, A., & Gormican, S. (1988). Feature analysis in early vision: Evidence from search asymmetries. *Psychological Review, 95*, 14–48.
- Treisman, A., & Sato, S. (1990). Conjunction search revisited. *Journal of Experimental Psychology: Human Perception and Performance, 16*, 459–478.
- Vitkovitch, M., & Humphreys, G.W. (1991). Perseverant responding in speeded picture naming: It's in the inks. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 17*, 664–680.
- Vitkovitch, M., Humphreys, G.W., & Lloyd-Jones, T.J. (1993). On naming a giraffe a zebra: Picture naming errors across different object categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 19*, 234–259.
- 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.
- Young, A.W., Newcombe, F., Hellawell, D., & De Haan, E. (1989). Implicit access to semantic information. *Brain and Cognition, 11*, 186–209.

*Manuscript received 1 August 1994*

*Revised manuscript received 1 June 1995*

## APPENDIX

Through a careful examination of the identification errors for ELM as well as several subjects (using multi-dimensional analysis), we hypothesized three simple shape parameters that are implicated in recognition; the three shape parameters we consider can be described as global transformations of a circle situated on the two-dimensional plane. Thus, they can be conceptualized as a set of three independent functions that are applied sequentially to an input circle, which deform it into one of a continuous family of output shapes. These output shapes are thus defined or represented by the set of three deformation parameters (one for each function). The advantage of this is, of course, that shapes can be readily compared to one another in terms of the explicitly defined parameter space. Note, however, that this representation cannot produce all possible two-dimensional shapes.

Starting with a unit circle, the specific parametric shape transformations we use are the following:

- change of elongation (i.e. stretching or the alteration of the ratio of minor to major axes for a defining ellipsoid),
- tapering along the major axis of the ellipsoid (i.e. making a cylinder-like object appear cone-like),
- and global curvature (blending) of the ellipsoid perpendicular to the major axis.

These were defined using sequential global transformations of the coordinate plane that map points  $(x, y)$  into new transformed points  $(x, y)'$ . These can be described as a combination of a rotation and translation about the origin for each point on the plane.

$$(x, y)' = \vec{p}' = \hat{\Omega}(\vec{p})\vec{p} + \hat{T}(\vec{p}) \quad (1)$$

where  $\vec{p}$  and  $\vec{p}'$  are the initial and transformed points, and  $\hat{\Omega}()$  and  $\hat{T}()$  are the rotation and translation matrices. The transformations we use are specified formally as follows (deformations are given with respect to the  $y$ -axis, with full generality).

To set the eccentricity of the shape to the value  $\alpha$ , we apply the transformation given by:

$$\hat{\Omega}(\vec{p}) = \begin{bmatrix} 1 & 0 \\ 0 & \alpha \end{bmatrix}, \quad \hat{T}(\vec{p}) = \{0, 0\} \quad (2)$$

This simply stretches points upwards along the plane by a factor of  $\alpha$ .

For tapering by a factor of  $\gamma$  we have:

$$\hat{\Omega}(\vec{p}) = \begin{bmatrix} \gamma/(K-y) & 0 \\ 0 & 1 \end{bmatrix}, \quad \hat{T}(\vec{p}) = \{0, 0\} \quad (3)$$

where  $K$  is a constant larger than the maximum  $y$  value of any shape of interest. This corresponds to a contraction in towards the  $y$ -axis of the plane by an amount that depends inversely on distance from a selected "focal point".

For bending by a factor  $\beta$ , we have:

$$\hat{\Omega}(\vec{p}) = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad T = \{0, \beta |y|\} \quad (4)$$

In practice, the thickness (and hence the area) of the generated shapes is held constant over when samples are generated. This is achieved by adjusting the eccentricity parameter to compensate for changes of thickness which arise from changes in tapering. As such, the actual object elongation (ratio of height to maximum width) is *not* identical to the eccentricity value  $\alpha$  except for objects with zero tapering.

Some sample shapes generated by this method, starting from an initial circle, are shown in Fig. 4.





Copyright of Visual Cognition is the property of Psychology Press (T&F) and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.