Semantic and Visual Determinants of Face Recognition in a Prosopagnosic Patient

Mike J. Dixon
University of Waterloo, Canada

Daniel N. Bub
University of Victoria, Canada

Martin Arguin
Université de Montréal, Canada.

Abstract

■ Prosopagnosia is the neuropathological inability to recognize familiar people by their faces. It can occur in isolation or can coincide with recognition deficits for other nonface objects. Often, patients whose prosopagnosia is accompanied by object recognition difficulties have more trouble identifying certain categories of objects relative to others. In previous research, we demonstrated that objects that shared multiple visual features and were semantically close posed severe recognition difficulties for a patient with temporal lobe damage. We now demonstrate that this patient's face recognition is constrained by these same parameters. The prosopagnosic patient ELM had difficulties pairing faces to names when the faces shared visual features and the names were semantically related (e.g., Tonya Harding, Nancy Kerrigan, and Josée Chouinard—three ice skaters). He made tenfold fewer errors when the exact same faces were associated with semantically unrelated people (e.g., singer Celine Dion, actress Betty Grable, and First Lady Hillary Clinton). We conclude that prosopagnosia and co-occurring category-specific recognition problems both stem from difficulties disambiguating the stored representations of objects that share multiple visual features and refer to semantically close identities or concepts. ■

INTRODUCTION

Prosopagnosia is the neuropathological inability to recognize familiar people by their faces. The temporal lobe patient ELM, for example, is unable to recognize the faces of his wife, sons, or grandchildren. He claimed a picture of the first author was unfamiliar, despite sitting immediately beside him.

Prosopagnosic patients can have intact perception. ELM can copy complex figures and animals. He can name photographs of objects taken from both standard and unusual views. He can match standard and unusual views of animals and artifacts and can select a target face from an array of distractor faces surrounding the target. Despite these intact perceptual abilities, over several years of testing ELM has never once spontaneously identified a single face. He cannot discriminate familiar from unfamiliar faces or previously viewed faces from novel unfamiliar ones, nor can he identify emotional expressions. (A more complete case report is given in the "Subjects" and "Methods" section following the discussion.)

Prosopagnosia patients can vary in the severity of their face recognition deficits. The prosopagnosia patient PV (Sergent & Poncet, 1990) was tested on a two alternative, forced-choice, face-name-matching test in which a face was presented along with the correct name and another name referring to a person of the same gender and occupation as the correct alternative. PV selected the correct name on 40 out of 48 trials. By contrast, ELM performed at chance levels on this same task (20/48 trials correct).

Prosopagnosic patients can also vary as to whether or not they display covert face recognition. Covert recognition refers to the fact that although patients cannot spontaneously name a face or pick the right name for a face from a set of names, using certain indirect behavioral measures (evoked potentials, galvanic skin response, semantic priming, and learning of true and untrue face-name pairings), some patients do show a certain degree of intact recognition (see Young, 1994, for a review). In a covert recognition paradigm using semantic priming Young, Hellawell, and De Haan (1988) asked the patient PH to make speeded familiarity judgments to printed names (e.g., John Lennon). The name was preceded by a related face (e.g., Paul McCartney), an unrelated face (Ronald Reagan), or a neutral unfamiliar face. Despite PH's inability to overtly identify these faces, his reaction times for the related condition (1016 msec) were significantly faster than the neutral (1080 msec) and unrelated conditions (1117 msec).

On a variant of Young et al.'s (1988) semantic priming task (using the related and unrelated but not the neutral condition), ELM showed no priming, and was actually marginally faster at saying a name was familiar when an unrelated face preceded the name (1037 msec unrelated versus 1056 msec related). A similar lack of priming effects was shown by Newcombe, Young, and De Haan's (1989) patient MS. Patients like ELM and MS, who fail to demonstrate covert face recognition, have been characterized as having a more severe degree of prosopagnosia than other patients who show covert recognition of faces (Farah, O'Reilly, & Vecera, 1993).

Although patients like PH show that intact semantics can lead to covert recognition, Sergent and Poncet's (1990) patient PV shows how semantics can occasionally allow covert recognition to become overt. This patient demonstrated a large degree of covert face recognition across a number of paradigms. When the faces of eight politicians were presented as a group, and PV was told that they were from the same occupational category, PV spontaneously named the category and then proceeded to name seven of the faces and gave accurate semantic information concerning the eighth. Similarly, with a set of eight actor's faces, she spontaneously named the category and correctly named all eight exemplars. Importantly, although PV could identify faces when they formed groups of actors or politicians, these same faces could not be identified when the actors and politicians were intermixed and presented individually. Sergent and Poncet suggested that "neither the facial representations nor the semantic information were critically disturbed in PV and her prosopagnosia may thus reflect faulty connections between faces and their memories" (Sergent & Poncet, p. 1000). They surmised that by presenting several members of the same semantic category together, they may have temporarily raised the signal across these faulty connections and surpassed the thresholds necessary for identifying both the category of the presented group and its individual members.

Although prosopagnosia can occur on its own, it is often accompanied by object recognition difficulties. The co-occurance of face and object recognition deficits may be indicative of a more general deficit that occurs "earlier" in the object recognition sequence. Such early problems may disrupt the recognition of both faces and objects and may preclude even the covert recognition of faces. ELM is a patient who presents with both profound prosopagnosia and object recognition deficits. Importantly, however, his recognition deficits are only for certain categories of objects. In confrontation naming ELM identified 92% of line drawings depicting artificial objects but only 21% of those depicting biological objects. Closer investigation revealed that the biological versus artificial distinction was itself a consequence of ELM's propensity to confuse objects that were visually similar and semantically proximate (e.g., fruits and vegetables, animals and birds, but also, cars and stringed musical instruments).

In investigating ELM's prosopagnosia, we were guided by the same structural and semantic factors that seemed to constrain his object recognition. Concerning structure, Arguin, Bub, and Dudek (1996) previously showed that ELM was relatively good at identifying shapes when exemplars within a set could be discriminated by attending to a single visual feature such as elongation. Conversely, he showed marked impairments when shapes within a set shared multiple visual features (elongation and tapering). Examples of these single-dimension and conjunction shape sets are shown in Figure 1. Dixon, Bub, and Arguin, (1998) used the paradigm shown in Figure 1 to demonstrate that these structural impairments interacted with semantics in a remarkable way. On learning trials, shapes comprising either the single-dimension or the conjunction sets were paired with familiar sounds. On test trials the shapes appeared alone, and ELM had to remember the sound that the shape was paired with during learning trials. A series of eight learning trials were followed by eight test trials, with this pattern repeated until 96 learning and test trials had been presented in block 1 and repeated in block 2.

The sounds that were paired with shapes were either semantically similar (e.g., sound of a robin, crow, owl) or semantically disparate (e.g., sound of saw, helicopter, photocopier). ELM performed equally well when single-dimension sets were paired to semantically close or disparate sounds. For conjunction sets paired to closely related sounds ELM's shape identification was poor. When the exact same shapes were paired with unrelated sounds, ELM's identification performance, after an initial learning period, was flawless. This finding was replicated using verbal labels as well as digitized sound recordings. When semantically close verbal labels like "Mustang," "Trans-am," "Camaro," and "Corvette" were used, substantial recognition deficits ensued (67% total errors). His performance was vastly improved when similar shapes were paired with "lion," "wasp," "frog," and "hummingbird" (2.6% total errors). In total, ELM was tested on 16 such quadruplets. We gathered semantic similarity ratings from 30 normal control subjects for each of the 16 quadruplets. For conjunction sets (n = 16), ELM's block 2 errors were significantly correlated with normal's ratings of the quadruplets semantic proximities (r = 0.84, p < 0.01). In contrast, ELM's errors on single-dimension sets (n = 16) were uncorrelated with normal's ratings of semantic proximity (r = 0.06, *ns*).

This pattern of increased identification errors for sets of shapes that share multiple visual features and refer to semantically close concepts has since been replicated in another patient with category-specific visual agnosia. (Arguin, Bub, Dixon, Caille, & Fontaine, 1996).

These results can be interpreted within the framework of exemplar models that employ the notion of psychological distance. In such models objects are represented as points in multidimensional psychological space, and the smaller the distance between objects, the greater their propensity to become confused in memory (Estes, 1994). We proposed that ELM's deficit involved abnormal difficulties in disambiguating the repre-

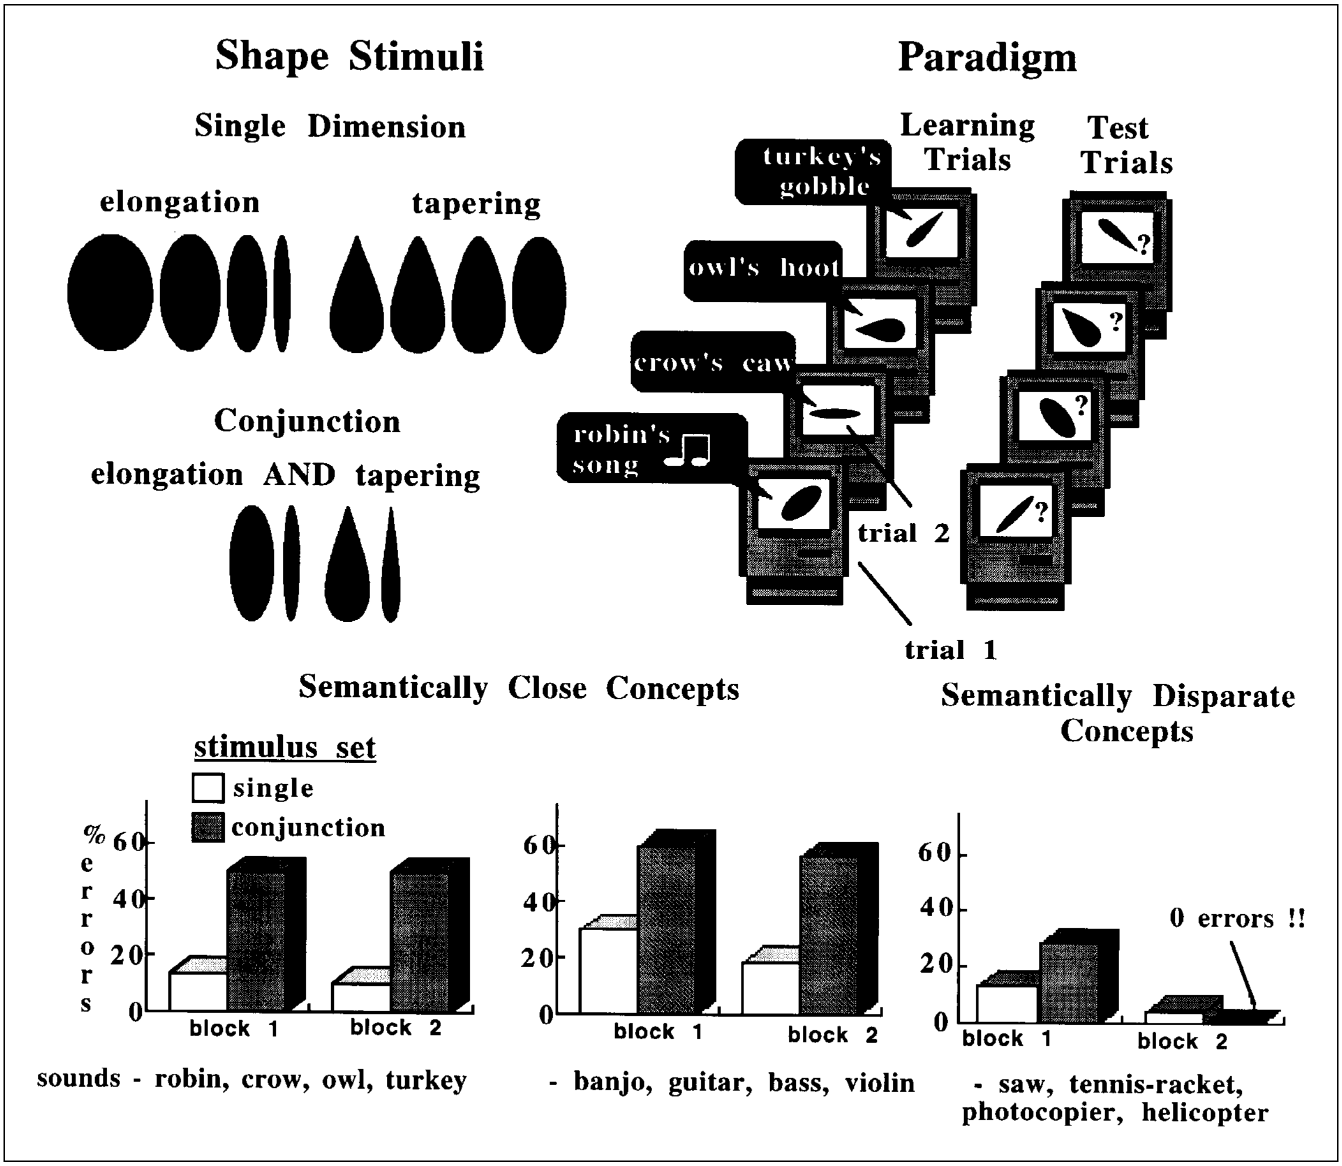

Figure 1. The computer-generated shape sets, paradigm, and results of Dixon, Bub, and Arguin, 1998. Results are the percentage of test trial errors for single dimension and conjunction shape sets associated with semantically close and semantically disparate concepts. Each block contains 96 test trials.

sentations of closely stored objects. The findings depicted in Figure 1 indicate that psychological distance depends on both structural and semantic factors. Associating shapes with semantically disparate labels appears to increase the psychological distance between exemplars and reduce confusions even when shapes share many visual features. When psychological distance is decreased by virtue of exemplars having both overlapping visual features and overlapping semantic attributes, identification deficits ensue.

We postulated that this combination of overlapping visual and semantic features might also be a contributing factor to ELM's prosopagnosia. We reasoned that although many faces share multiple visual features (e.g., eye, nose, and mouth size) within this category, faces whose identities also share many semantic attributes might be even more confusable than faces mapped to identities having few overlapping attributes. If so, the faces of hockey players, for example, should be more confusable than the faces of men having different professions.

In order to see if ELM's face identification was constrained by the same factors as his object recognition, we asked ELM to remember face-name pairings for two quadruplets of male model's faces. We used the same paradigm that was used for the shape-naming task depicted in Figure 1. On learning trials, faces were presented along with a digitized recording of a name. On test trials the faces were presented alone, and ELM had to "name" the face. Eight learning trials were followed by eight test trials, with this pattern repeating until 96 learning and test trials per block had been given. (Pre-

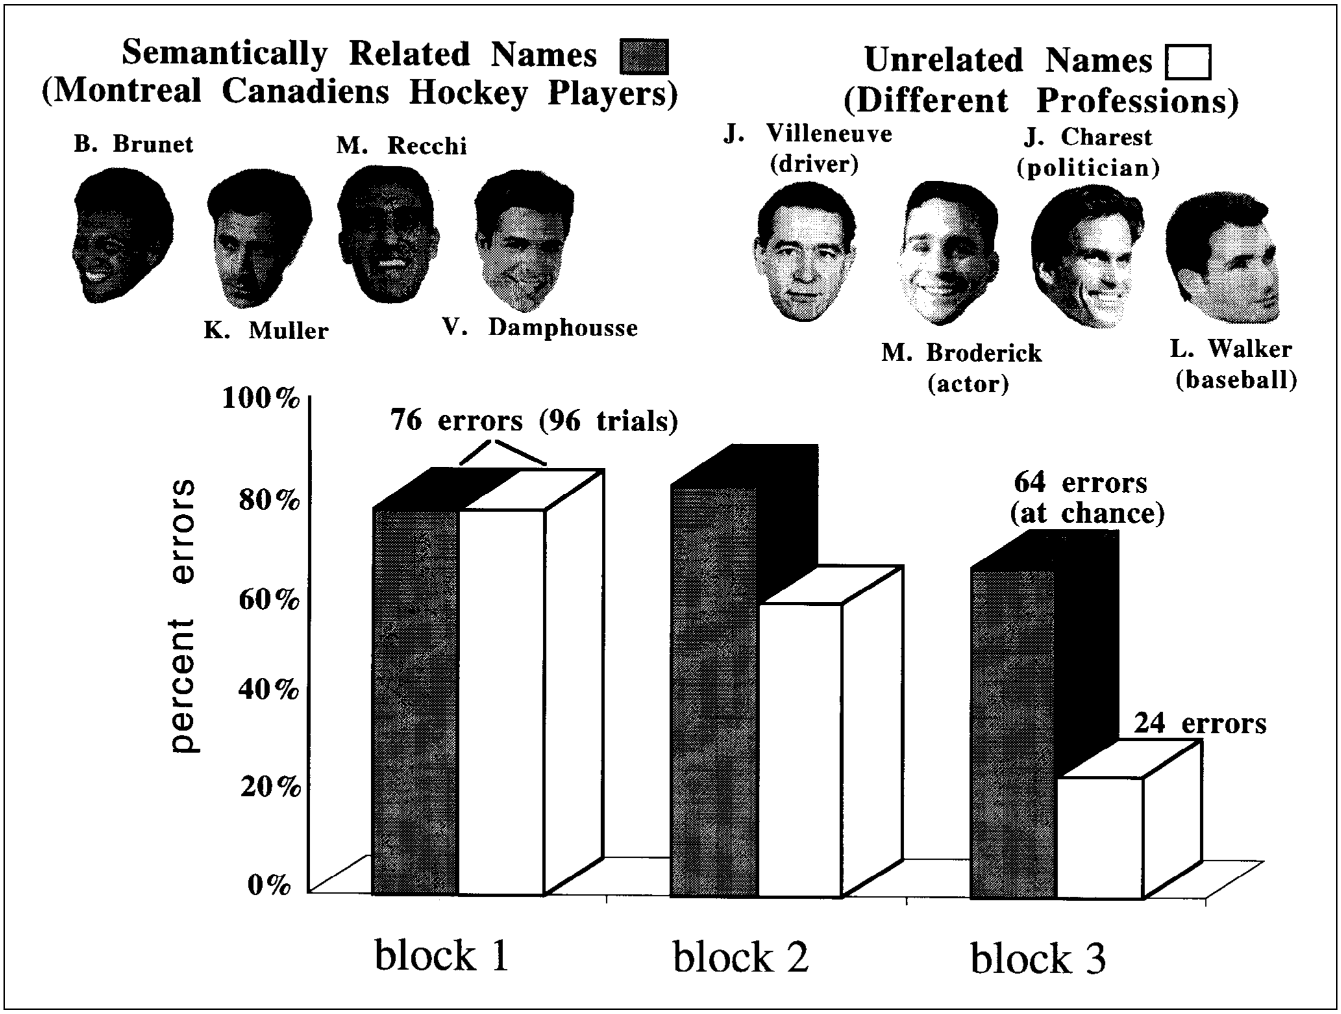

Figure 2. Percentage of test trial errors for sets of faces paired to semantically close and disparate names. Each block contains 96 test trials.

vious testing with blobs indicated that this interleaving of learning and test trials was necessary for ELM to learn to recognize even limited sets of stimuli.) Three identical blocks were administered per face set. For one set of faces the names used on learning trials were those of forwards on the Montréal Canadiens Hockey Club (a semantically similar set). In the other set the names paired with the faces were from different occupational categories: a race-car driver, a politician, a baseball player, and an actor (a semantically disparate set).

EXPERIMENT 1

Results

The results depicted in Figure 2 convey the severity of ELM's prosopagnosia. After 288 repetitions (blocks 1, 2, and 3), ELM was still at chance at identifying four faces paired with four semantically close names, and there was no change in performance over the three blocks [$\chi 2$ (2) = 1.88, *ns*]. For faces paired with semantically unrelated names performance improved over the three blocks, ($\chi 2(2) = 27.58$ $p < 0.001$). On block 3, performance was significantly better for the set using unrelated names than for the set using related names ($\chi 2(1) = 19.67$ $p < 0.001$).

EXPERIMENT 2

In interpreting ELM's *object* recognition deficits, we made the assumption that both ELM and healthy adults stored objects in memory according to the principle of psychological similarity. Accordingly, objects that have many overlapping visual and semantic features would be stored close together. Where ELM differs from normals is in the difficulty he has disambiguating such contiguously stored exemplars. Whereas normals only have difficulty disambiguating exemplars under unusual laboratory conditions such as time-limited naming (subjects may erroneously call a leopard "tiger"), ELM's lesion appears to cause difficulties disambiguating psychologically proximate exemplars in everyday life. Thus, he will confuse fruits and vegetables, animals, makes of car, and certain musical instruments.

Experiment 1 presents preliminary evidence that ELM's *face* recognition problems stem from this same inability to disambiguate psychologically close exemplars. Although faces in general prove too psychologically proximate for ELM to disambiguate in everyday life, Experiment 1 demonstrates that with limited sets of faces, and many learning trials, ELM can learn to identify certain faces. Faces that are stored relatively far apart in psychological space by virtue of having disparate semantics can be identified by ELM at levels in excess of chance. Faces that are psychologically very close to one another because of overlapping semantics are significantly more difficult for ELM to disambiguate, and under these conditions are identified only at chance levels.

Although the most parsimonious explanation of ELM's significantly poorer performance with the faces mapped to hockey player's names involves a detrimental effect of semantic proximity, it could be argued that ELM simply performed more poorly on this set because, by chance, these exemplars had more overlapping visual features than the set of faces mapped to semantically unrelated names. At present our knowledge of faces does not allow us to unequivocally state which key facial dimensions determine how similar one face is to another. Thus, although attempts were made to match face sets as closely as possible, we cannot rule out the possibility that faces within these sets differed in their visual proximity.

In order to rule out this explanation, and gain empirical control over the visual dimensions comprising each presented face, in Experiment 2 we employed a similar paradigm but used triads of *synthetic* faces. These faces were constructed using a template female face in which eye, nose, and mouth size could be varied. Using synthetic faces enabled us to eliminate nondiagnostic differences between faces, empirically specify exactly how faces within triads differed from one another, and force ELM to process multiple facial features for the purposes of identification. These faces are depicted in Figure 3.

Two sets of faces were used. One set was paired with semantically similar names (ice skaters—Tonya Harding, Nancy Kerrigan, and Josée Chouinard). A second set was paired with semantically disparate names (First Lady Hillary Clinton, actress Betty Grable, and singer Celine Dion). Each face set was then presented a second time using the opposite name set. To control for potential practice effects we used the ABBA design depicted in Figure 3.

Finally, ELM was asked to associate the three skaters' names to faces that were made visually distinct via distortion. This manipulation was used to show that ELM could learn to identify semantically similar entities provided they were psychologically distant because of their visual dissimilarity. Distorted rather than plausible faces were used because, given the severity of ELM's face identification difficulties, even the most discrepant plausible faces may have proved too visually similar for ELM to disambiguate when combined with semantically similar names. A good performance in this condition would refute the possibility that ELM had an impairment purely at the semantic level (associating semantically close labels to *anything*).

ELM's performance on each condition was compared to six healthy, independent living, age-matched controls without any subjective complaints concerning either memory or face recognition. In order to avoid ceiling effects among control subjects, six additional age-matched control subjects were tested using sets of four faces (two conjunction sets formed by combining two eye sizes with two mouth sizes) shown in Figure 3. These face quadruplets were associated with quadruplets of semantically similar or unrelated names (ice skater Katerina Witt and tennis player Steffi Graf were added to complete the relevant related and unrelated name sets).

Results

Figure 3 summarizes the design and performance of ELM and the second set of age-matched controls for all conditions.

When asked to learn face-name associations using names that were unrelated, ELM's identification performance was nearly flawless (only four errors on sets B_1 and B_2). His performance for the exact same faces paired with semantically related names was much poorer (60 errors on set A_1; 41 errors on set A_2). Using the most conservative comparison (A_2: 41 errors vs. B_1: 4 errors) ELM's performance was 10 times worse for faces that shared visual features and had semantically similar names than for the exact same faces paired with semantically unrelated names ($\chi 2(1) = 21.13, p < 0.001$). Performance was also 10 times worse for the visually similar faces paired with semantically close names than for faces in which the same names were paired with visually distinct faces (B_2 vs. Distorted Faces, $\chi 2(1) = 21.13, p < 0.001$).

ELM's confusions between faces in the semantically close sets were not arbitrary. Rather, they depended on the features that pairs of faces had in common. As can be seen in Figure 3, in set A_1 all three faces had the same nose, but the members of the triad differed in eye and mouth size. In this set 13/60 confusions were made between pairs of faces with the same sized eyes (faces labeled Tonya Harding and Josée Chouinard) and 41/60 confusions were made between the pair of faces with the same sized mouths (Chouinard and Kerrigan). Only 6/60 errors were made between the pair of faces that had different eye and mouth sizes. This pattern was replicated in set A_2 with 38/41 confusions made between faces with the same sized mouths and 3/41 confusions made between faces with the same sized eyes.

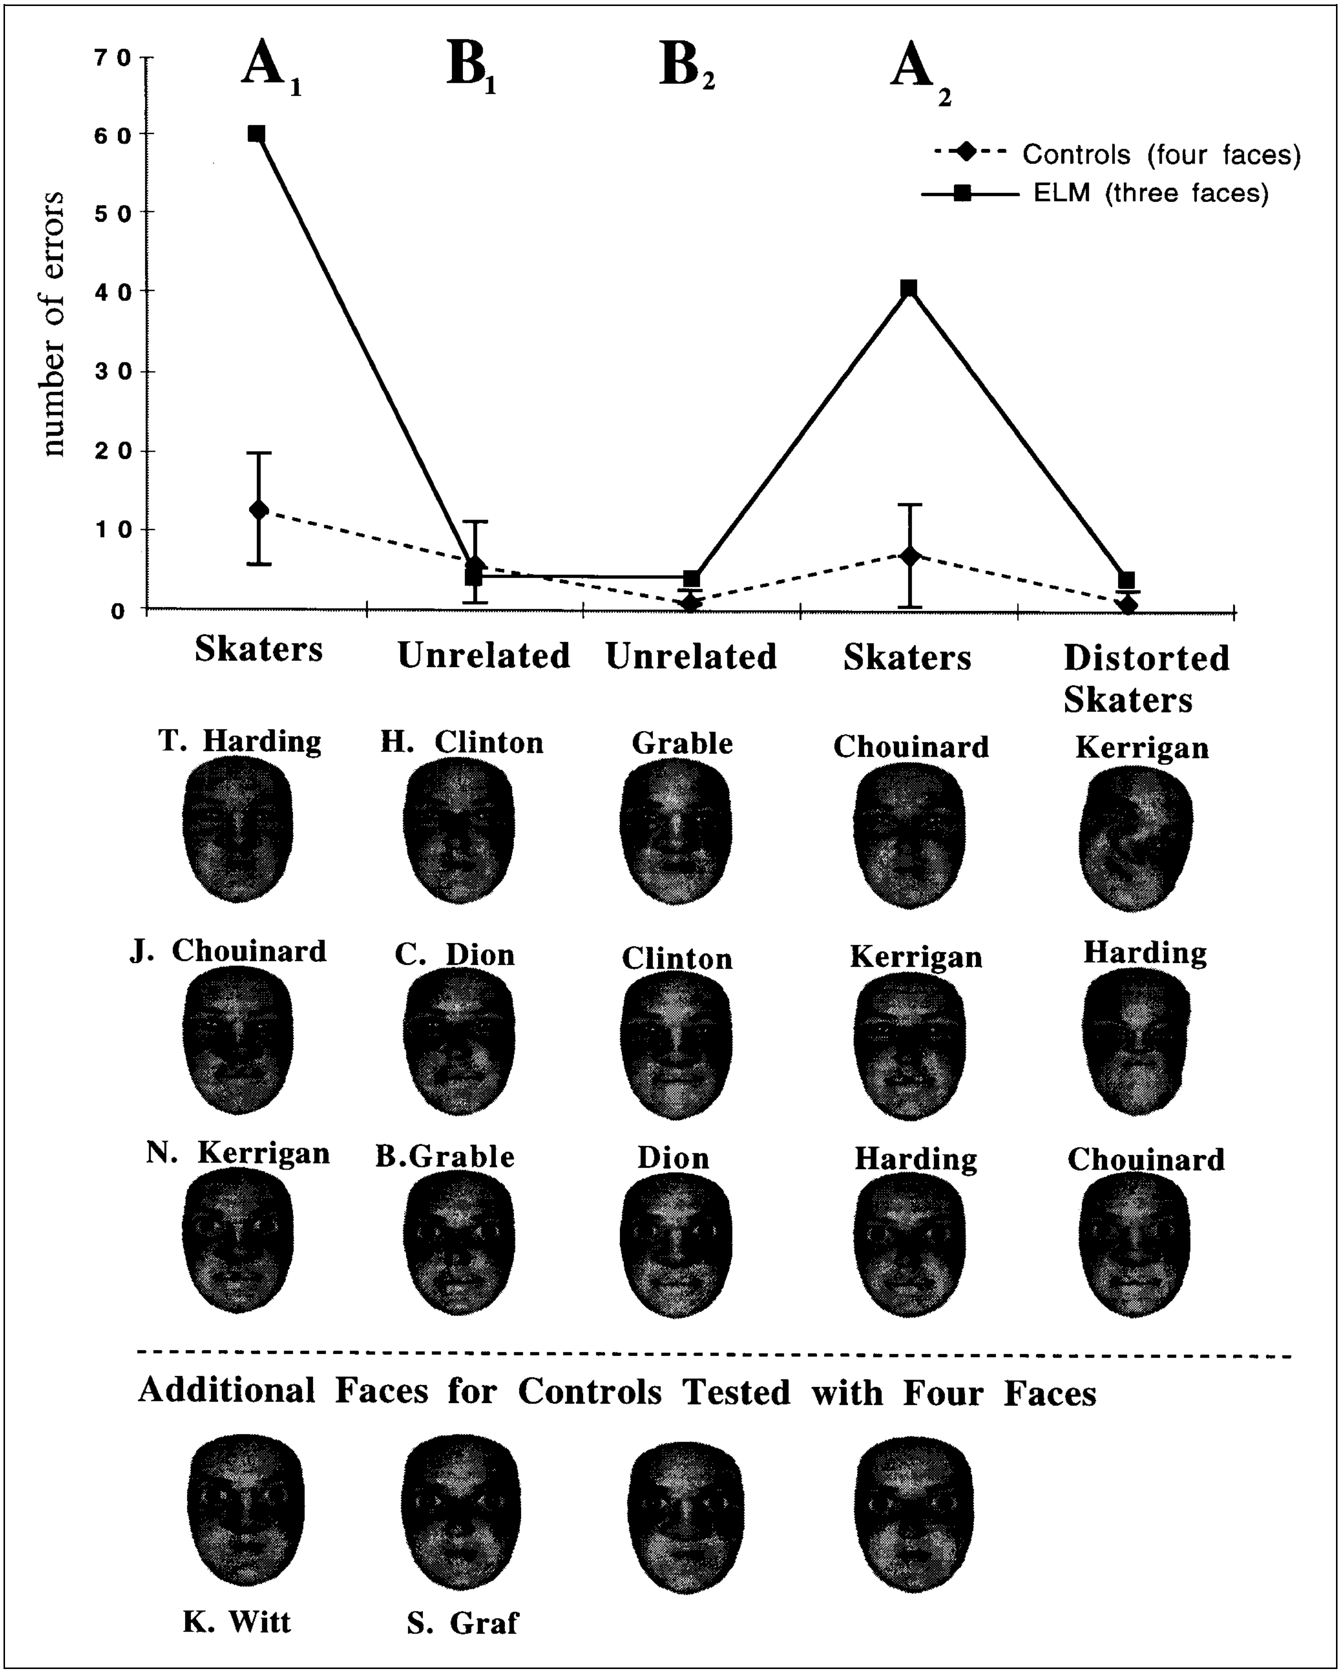

Figure 3. Test trial errors for faces associated with semantically close and disparate sets of names. The triads of faces and names used to test ELM and one set of controls are presented above the dotted line. Below the dotted line are the extra faces used for controls tested on quadruplets of faces. The face-name pairings for controls working with quadruplets of faces are those shown in A_1 and B_1. For sets B_2 and A_2 face-name pairings were assigned so that there was no overlap between A_1-A_2 and B_1-B_2 face-name assignments.

Table 1. Mean number of errors for triads of faces committed by ELM and the two sets of six aged-matched controls on the ABBA and distorted conditions. (Standard deviations for controls are in brackets.)

	A_1 *skaters*	B_1 *unrelated*	B_2 *unrelated*	A_2 *skaters*	*Distorted skaters*
ELM (three faces)	60.0	4.0	4.0	41.0	4.0
Controls (n = 6) (three faces)	4.67 (4.13)	1.83 (2.99)	5.17 (5.07)	0.67 (1.21)	0.33 (0.81)
Controls (n = 6) (four faces)	12.67 (8.48)	5.83 (6.34)	0.67 (1.63)	7.16 (7.44)	0.83 (1.60)

Faces that differed on both eye and mouth sizes (Harding and Chouinard) were never confused in this set. This pattern of confusions is exactly the same as the pattern that was obtained when ELM was asked to associate conjunction sets of blobs to semantically related (e.g., bird) names. Blobs differing in two dimensions were rarely confused, but many confusions occurred between blobs that overlapped on one of the two dimensions (Arguin et al., 1996).

ELM (Triads of Faces) versus Aged-Matched Controls (Triads of Faces)

The performance of ELM and controls working with triads and quadruplets of faces is shown in Table 1. For the controls working with triads of faces, 43% of errors were made among the semantically related sets and 47% among the unrelated sets. In contrast ELM made 93% of his errors in the semantically related sets and only 7% errors among the unrelated sets. These percentages are significantly different from the control percentages associated with the same conditions ($\chi 2(1) = 137.38, p < 0.001$)

In terms of the ABBA design, performance of normal controls showed neither order ($F(1,5) = 0.06$, *ns*) nor semantics ($F(1, 5) = 0.37$, *ns*) nor significant interaction effects. Performance was close to ceiling in all conditions, however, possibly precluding a semantic effect.

ELM's identification performance scores were converted to z scores using the means and standard deviations of the age-matched controls. ELM's performance was considered significantly different from controls if it lay three standard deviations beyond control means (statistical outliers are traditionally associated with z values of greater than 3.0, $p < 0.0013$). Using this criterion, ELM showed performance that was indistinguishable from controls on both triads of unrelated faces ($z = 0.72$ and $z = -0.23$ for B_1 and B_2, respectively). ELM performed significantly worse than controls on both triads of semantically related faces ($z = 13.40$ and $z = 33.30$ for A_1 and A_2, respectively). ELM also performed significantly worse than controls on the visually distorted face triad ($z = 4.49$).

ELM (Triads of Faces) versus Aged-Matched Controls (Quadruplets of Faces)

For controls working with quadruplets of faces, overall, 75% of errors were made among the semantically related sets and 25% among the unrelated sets. ELM's rates of 93% for semantically related sets and only 7% errors for unrelated sets were significantly different from the control percentages associated with semantically related and unrelated conditions ($\chi 2(1) = 12.05, p < 0.01$).

Control subjects working with quadruplets of faces performed significantly better on the unrelated conditions relative to the semantically related ones of the ABBA design ($F(1, 5) = 9.08, p < 0.05$). There was no effect of order ($F(1, 5) = 4.94$, *ns*) for these six control subjects.[1]

Using the z-score transformations, ELM again showed performance that was indistinguishable from controls on both sets of unrelated faces ($z = -0.29$ and $z = 2.04$ for B_1 and B_2, respectively). ELM performed significantly worse than controls on both sets of semantically related faces ($z = 5.58$ and $z = 4.41$ for A_1 and A_2, respectively). ELM's error rates were comparable to those of controls in the distorted condition ($z = 1.97$).

Discussion

The combined results of Experiments 1 and 2 illustrate a profound effect of semantics on ELM's ability to identify faces. When faces are visually *and* semantically similar, severe recognition problems ensue. ELM's problem is not purely at the level of semantics. That is, he does not merely have a problem disambiguating exemplars that are semantically related. If this were the case, ELM would have had trouble with the visually distorted faces mapped to the semantically related ice skaters' names. Consonant with the psychological similarity storage principle, by making these faces more visually distinct via distortion, identification problems were dramatically reduced. Thus, ELM's problem seems to involve the interaction of structural and semantic factors. When exemplars within a set are either visually distinct or semantically unrelated, ELM can learn to identify limited numbers of faces. If, however, exemplars within a set are both visually and semantically similar, profound recognition problems are encountered.[2]

This paradigm used in Experiment 2 is a novel approach to decoupling structural and semantic factors in face recognition. By associating the exact same faces to semantically close and disparate names in an ABBA design, one can in essence hold the structural factors constant and look directly at the profound effect of semantics on face recognition. Comparing the findings illustrated in Figure 1 with those in Figures 2 and 3, it is apparent that the same structural and semantic factors constrained both ELM's face and object recognition.

GENERAL DISCUSSION

The finding that semantic proximity affects face identification in both ELM and healthy control subjects has important ramifications for models of face and object recognition. Currently, the approaches used to train computers to recognize faces or objects focus only on visual features (e.g., Beymer & Poggio, 1996). The performance of ELM and controls in Experiment 2 indicates that in humans, semantics also plays a key role in face recognition.

ELM's propensity to confuse faces and objects that are both visually and semantically similar is reminiscent of the "mixed" visual and semantic reading errors (e.g., reading "rat" as "cat") made by patients with deep dyslexia (Hinton & Shallice, 1991). In simulations of many (but not all) connectionist configurations, mixed errors are a prevalent form of deep dyslexic reading errors and can arise from damage to a number of different model components (Hinton & Shallice, 1990; Plaut & Shallice, 1993).

For object recognition, there may be an even greater propensity to confuse items that are visually and semantically related. This is because in reading, word form is only arbitrarily related to semantics (the word "cat" does not look like the four-legged feline), and large attractor basins are proposed to overcome the problems distributed architectures have in making visually similar inputs (cat, mat) elicit disparate patterns of activations in semantic space. For visually presented objects, on the other hand, visually similar forms are often also semantically similar (e.g., the shapes of a robin and a crow are similar, as are their meanings). Distributed architectures require less learning and smaller connection strengths to have visually similar inputs elicit similar but still discernibly different patterns of activation in semantic space (Plaut & Shallice, 1993). As in deep dyslexia, however, a consequence of storing object representations in this distributed fashion would be that damage to this architecture would likely result in the propensity to confuse objects that are both visually and semantically related. The same would apply to faces stored in distributed fashion; damage would preferentially disrupt the recognition of visually and semantically similar faces. Thus, whether accounting for the word reading errors of deep dyslexic patients or the face and object recognition problems of temporal lobe patients like ELM, if one assumes that knowledge is stored in a distributed architecture, confusions among entities that are both visually and semantically similar are to be expected when this architecture becomes damaged.

One nondistributed model of face recognition that includes both visual and semantic components is that of Burton, Bruce, and Johnston (1990). ELM's preferential difficulty identifying semantically similar faces poses an interesting challenge for this model. The model uses a connectionist architecture to link different pools of units devoted to encoding specific processes relevant to face recognition. It consists of pools of face recognition units (FRUs), which are connected to person identity nodes (PINs), which, in turn, are connected to semantic information nodes (SINs). Importantly, there are no direct connections between the FRUs and the SINs. Nodes within a pool are negatively connected (inhibit one another), whereas nodes in different pools are positively connected and can facilitate one another bidirectionally. The visual aspects of an individual face (e.g., that of Prince Charles) is coded by a specific FRU, which then facilitates a specific unit in the person identity node pool. A specific PIN will then excite specific units in the semantic information pool relevant to that particular person. For example, the Prince Charles PIN will facilitate a "royalty" node in the semantic pool. This royalty node, in turn, elevates the activation levels of PINs to which it is directly connected (e.g., the Prince Charles PIN but also the Lady Diana PIN). This model can account for semantic face priming, identity priming, cross-modal priming (face-name priming in a familiarity decision task) and the fact that visually distinct faces can be more rapidly identified than faces with more typical features.

In order to account for ELM's preferential difficulties with semantically similar sets of faces, one could assume that ELM's lesion has reduced the negative, inhibitory connections between different person identity nodes. Thus, in Experiment 2 it could be proposed that the PIN associated with Hillary Clinton would activate certain semantic nodes (e.g., a "First Lady" node) that would project back to the Hillary Clinton PIN and increase its activation. Lacking direct connections, the First Lady semantic node would not facilitate the Betty Grable or Celine Dion PINs. Thus, relative to the other two PINs in the triad, the Hillary Clinton PIN would be much more activated, and hence, easier to disambiguate from the less-activated Grable and Dion PINs.

In the semantically related condition, when the Josée Chouinard PIN is activated, the semantic information nodes ("ice skater," "Olympics," etc.) would facilitate the PINs for Chouinard, Kerrigan, and Harding equally, perhaps making them harder to disambiguate because of a failure of the inhibitory connections to drive down semantically related, but incorrect, PINs. If this interpretation is correct, the challenge to the Burton et al. model then becomes how to account for what is in essence the opposite finding, namely, that presenting sets of semantically close faces to Sergent and Poncet's (1990) patient PV facilitated rather than impaired face identification.

Besides the obvious paradigmatic discrepancies, the different effect of semantics may be attributable to differing severity levels of prosopagnosia. PV demonstrated good performance on a two-alternative, forced-choice, face-name matching task (40/48 correct), whereas ELM was at chance on this task (20/48). PV showed covert recognition for faces in a number of paradigms, whereas ELM did not. Recall that when shown a face (e.g., Prince Charles) and then asked to make a speeded familiarity judgment concerning a subsequent name (e.g., "Diana Spencer"), ELM showed no priming and was marginally faster when the face was *unrelated* to the name (1037 msec unrelated versus 1056 msec related). Recall also that patients who fail to demonstrate covert recognition have been interpreted as having more severe face recognition deficits than patients who show covert recognition (Farah et al., 1993).

Thus, in defense of the Burton et al. (1990) model, the opposing costs and benefits of semantic proximity for ELM and PV could be attributable to different levels of severity. How then would the model fare in comparing ELM to a patient with prosopagnosia of comparable severity? Like ELM, Newcombe et al.'s (1989) patient MS failed to show priming in the face-name priming task. Also like ELM, this patients' severe prosopagnosia was accompanied by object agnosia. Together, these findings suggest that ELM and MS suffer from prosopagnosia of comparable severity. Within the Burton et al. model, MS was interpreted as having higher-order perceptual impairments that prevented adequate information from entering into the FRUs. Despite having prosopagnosia at least as severe as that of MS, the current study shows that semantic proximity profoundly affected ELM's ability to disambiguate one face from another in memory. This suggests that semantics can have profound effects even at relatively "early" stages in the face recognition sequence. Placing ELM's deficit prior to the FRUs (like MS) or alternatively at the level of the FRUs may be problematic for the model of Burton et al. because there are no direct connections between FRUs and semantics. If ELM's deficit is interpreted as affecting disambiguation at the level of the person identity node pool (rather than at the face recognition node pool), one must account for why he fails to show priming in the face-name priming task and why patient PV shows enhanced rather than exacerbated performance with semantically similar faces.

It may be possible to explain the performance of ELM and PV within the framework of some exemplar models of category learning (Estes, 1994; Kruschke, 1992; Nosofsky, 1986). In such models, objects are represented as arrays of attributes. For our purposes we will apply this model to the representation of faces. Although individual faces can be distinguished using a myriad of visual features, for ease of explication our discussion will limit the number of facial dimensions to only the two diagnostic features used in Experiment 2: eye size and mouth size. Thus the faces in set A_1 in Figure 3 could be coded as (1,1) for the face labeled Tonya Harding, (1,2) for Josée Chouinard, and (2,2), for Nancy Kerrigan. In this coding scheme the value 1 represents small eye or mouth sizes, and 2 represents large eye or mouth sizes.

In Kruschke's (1992) ALCOVE model, there are three essential layers: input nodes, hidden exemplar nodes, and output nodes. Each input node encodes stimulus values on a single dimension (for our purposes the relevant dimensions are eye size and mouth size). The hidden exemplars nodes, to which inputs are connected, are represented as points in a multidimensional psychological space. These points are located using coordinates that are based on the dimension values that comprise a given exemplar (coordinates 1,1; 1,2; and 2,2, above). Hidden exemplar nodes have activation profiles. They respond most strongly to stimuli that have the same values as their location coordinates (e.g., the exemplar node located at 2,2 would respond most strongly to the face with large eyes and a large mouth but still would respond to exemplars comprised of similar values (e.g., 2.1, 1.9). Activation falls off exponentially as similarity between the input stimulus and the exemplar decreases. The spatial extent of these activation profiles, which Kruschke refers to as the hidden nodes "receptive field," depends on a specificity parameter. Large specificities mean that hidden nodes will respond only to stimuli very close to the exemplars that they code. That is, the (2,2) face would respond to a face with eyes and mouth values of (2.1, 1.9) but not to a face with values of (1.5, 1.5). Small specificities yield larger receptive fields, meaning that nodes will still respond even to stimuli that only loosely resemble the exemplars that they represent (e.g., exemplar node located at position (2,2) would still become activated by face (2, 1.5) and even (1.5, 1.5) but not (1, 1). Output nodes representing categories have learned connection strengths to the hidden exemplar nodes.

ELM's face identification performance can be interpreted most parsimoniously by assuming a deficit in the specificity parameter. That is, ELM would have small specificities yielding much larger receptive fields than normal—so large that the receptive fields for different exemplars overlap. Thus for ELM the hidden exemplar that responds maximally to the face labeled Nancy Kerrigan (2,2) would also respond to faces having values of (2,1)—the face labeled Tonya Harding—and also to the face having values of (1,2)—the face labeled Josée Chouinard. Empirically, 90% of his confusions on set A_1

(54/60 errors) and 100% of confusions on set A_2 were between faces like these that shared the same values on either eye size or mouth size.

Although wider than normal, ELM's confusions indicate that the receptive fields for his hidden exemplar units are not infinitely wide. That is, he will seldom confuse faces that do not share diagnostic features [e.g., Harding's face (exemplar node (1,1) in set A_1) with Kerrigan's face (exemplar node (2,2) in set A_1)]. Empirically only 10% (6/61) of his errors were of this type for set A_1 and 0% (0/41) for set A_2.

ELM's vastly superior performance when faces are mapped to semantically unrelated identities is explained by the fact that exemplar similarities are not based on visual feature information alone. In Estes' (1994) exemplar model, the confusability of exemplars is captured by both the visual similarity parameter s and the semantic similarity parameter σ. Both parameters are calculated in the same way (using a product rule), and a total "psychological" similarity value is arrived at by multiplying the visual similarity parameter by the semantic similarity parameter. Thus, in Estes' model, because semantically close concepts share a large number of conceptual attributes, they would have large values of σ, whereas semantically distant concepts would have small values of σ. Because overall psychological similarity = $s\sigma$, theoretically two sets of objects (or in our case, faces) having the exact same visual attributes (e.g., eye, nose, and mouth size) could have very different psychological similarity values depending on the semantic proximity of their labels (σ). Such is obviously the case for the faces in Figure 3. For visually similar faces mapped to semantically related names, the overall psychological similarity is large, and confusions are rampant. When semantically disparate names are used to label these exact same faces, this lessens overall psychological similarity, and faces are much less confusable. In other words, if one thinks of faces as occupying points in multidimensional space, the location of these faces must be determined by both visual factors (what the face looks like) and semantic factors (to whom the face belongs).

For ELM, problems arise when he must extract from memory faces that are close together within this psychological space by virtue of exemplars sharing visual and semantic attributes.

Looking at the control subjects who mapped quadruplets of faces to quadruplets of names, one can see the ramifications of both the psychological similarity storage principle and the ramifications of ELM's abnormally wide hidden node receptive fields. Although not excessively difficult, mapping four faces to four names elicited a substantial number of errors among these healthy subjects. Consonant with the psychological similarity storage principle, significantly more errors were made for faces in the semantically related sets than in the psychologically more distant semantically unrelated sets. This shows that healthy adults, as well as prosopagnosics like ELM, store faces according to a psychological similarity principle.

The ramifications of ELM's abnormally wide receptive fields can be seen in his dramatic exacerbation of errors for the semantically related sets. Although ELM does not differ from controls in his performance on the semantically unrelated sets, for the semantically related sets his error rates were 5 times as great as the error rates associated with controls doing quadruplets of faces and 15 times as great as the error rates associated with controls doing triads of faces. If ELM has abnormally wide receptive fields for individual face exemplars, he will have the greatest problems disambiguating faces that are stored close together in multidimensional psychological space. Thus, imbuing faces with overlapping semantics causes these faces to be stored much closer together in psychological space—rendering disambiguation of these close exemplars extremely difficult because of these abnormally wide and overlapping receptive fields.

The behavior of PV can be interpreted within a similar framework. According to Sergent and Poncet (1990) both the face and semantic representations associated with faces are intact, but the connections between these two subsystems are faulty. Thus, unlike ELM, at the level of structural representation, PV can be assumed to have normal (i.e., small, nonoverlapping) receptive fields for individual exemplar nodes in the hidden unit layer. If, as the data from Experiments 1 and 2 suggest, the location of exemplar nodes in multidimensional psychological space is determined using both structural and semantic coordinates, then faces belonging to people with overlapping semantic attributes (e.g., politicians) would be located close together in this multidimensional space. Thus, when PV is shown eight faces all belonging to the same category, this would serve to activate eight exemplars in the same area of this multidimensional space. Because hidden nodes are connected to category nodes in ALCOVE, such simultaneous activation of exemplars connected to the same category node may have pushed this category node (e.g., the node corresponding to "politicians") above threshold. Once the category is known, PV could access her intact semantics concerning members of this category. From this point PV could exert a type of "positive top-down influence on perception, sharpening the processing of facial features" referred to by De Renzi, Faglioni, Grossi, and Nichelli (1991, p. 219), in order to successfully identify each individual face. Unlike ELM, she does not confuse exemplars within categories because she has normal-sized, nonoverlapping receptive fields for individual exemplars.

When patient PV is presented with individual faces of people from different categories (e.g., politicians, actors, etc.), this would cause points from different areas of the multidimensional space to become activated, leading to increases in the activation levels of multiple category nodes. Because of faulty connections, these individual increases in the activation levels of the multiple category nodes would remain well below threshold, and unique identification of exemplars belonging to these categories would be impossible.

Viewed in these terms, it is possible to show how different problems (PV's impaired connections from structure to semantics versus ELM's overlapping receptive fields) might cause the semantic proximity of identities to have opposing effects on face recognition. In summary, the interpretations above are not meant to constitute a formal model of face recognition, but rather they use modifications of existing exemplar models, and in particular the notion of "psychological similarity," as a framework for thinking about diverse prosopagnosic phenomena.

Irrespective of one's views concerning the utility of applying exemplar models to face recognition, the research presented in Experiments 1 and 2 emphatically suggests that semantic factors can strongly influence the ease with which faces can be identified by a severely prosopagnosic patient. This research suggests that ELM suffers from a general failure to disambiguate the stored representations of exemplars that share multiple visual features and refer to semantically close concepts. That both object and face recognition are constrained by the same factors provides fundamental clues concerning temporal lobe function.

It would seem that the temporal lobes (at least areas 21 and 37) are of key importance in retrieving from memory fully specified descriptions of object form (Arguin et al., 1996) and that damage to these structures has the potential to disrupt recognition when the form of objects share visual features. Importantly, this potential is only fully realized when, in addition to sharing visual features, objects also have overlapping semantic attributes (Dixon et al., 1998).

To account for the category-specific nature of ELM's recognition problems, it can be postulated that many artifacts belonging to the same semantic category (e.g., saw, pliers, and hammer) are visually dissimilar and hence pose no recognition problems. Furthermore, nonbiological objects may pose fewer problems than biological objects because they have specific and often unique functions (Warrington & McCarthy, 1983, 1987, 1994). Recent brain imaging evidence by Martin, Wiggs, Ungerleider, and Haxby (1996) indicated that identifying animals draws upon primary visual cortex and ventral areas of the temporal lobes (the area damaged in ELM). When subjects identify tools, however, temporal lobe activation is accompanied by activation in the frontal lobes. Martin et al. attributed this frontal activity as reflecting activation of the brain area responsible for encoding knowledge about object use. In exemplar model terms, if nonbiological objects such as tools have different functions, as well as different forms, this would serve to increase the semantic distance between the exemplars comprising nonbiological categories, thereby making these objects easier to recognize. Exceptions are categories such as makes of car and musical instruments, which have similar forms and similar functions. Importantly it is these categories of nonbiological objects that pose problems for ELM and other category-specific agnosics who otherwise have problems only with living things (Damasio, 1990).

Unlike most artifacts, biological objects not only share a large number of semantic features but also share a large number of visual features (all animals have heads, necks, trunks, and legs). Furthermore, they typically do not have unique functions or other distinctive attributes to aid in distinguishing among exemplars within categories. Thus, objects like fruits, vegetables, animals, birds, and insects pose the fatal combination of semantic proximity of concepts and shared values along critical shape dimensions that precludes object recognition in ELM.

Like animals, faces share a plethora of visual features (all faces have eyes, a nose, a mouth, chin, etc.) Faces pose an even more difficult recognition problem than most biological objects for three reasons. First, whereas many biological objects can be distinguished from related objects using the presence or absence of unique features, (tiger has stripes, lion does not have stripes), faces must be distinguished using relatively small differences in features that all members of the category possess. That is, one may differentiate two faces using differences in the size of a given feature (e.g., nose size) but not the absence of a feature. Second, whereas animals or fruits and vegetables are typically identified using basic-level categories (e.g., tiger or apple), faces must be identified at the level of the individual (John Smith, rather than human being). Finally, the number of exemplars in the category of human faces is huge compared to the numerosity of most other categories. This means that although one can demonstrate the profound effects of semantics on ELM's face recognition in the laboratory, in real life, the sheer number of faces that must be disambiguated for unique identification means that there will always be exemplars that share multiple visual features and have overlapping semantics. Unfortunately, for people with mesiotemporal lobe damage of the kind sustained by ELM, the problem of disambiguating such visually and semantically similar exemplars will prove insurmountable.

SUBJECTS and METHODS

Case Description

Clinical History

ELM is a 69-year-old retired man who formerly worked in the purchasing department of a manufacturing plant. In December of 1982 ELM was admitted to hospital for heart failure. Neurological symptoms of sudden onset were reported on December 5, 1982. These included nominal dysphasia, left/right confusion, dyscalculia, and agraphia without alexia. An emergency computerized tomography scan revealed a hypodensity deep in the right mesiotemporal lobe. The neurological symptoms resolved and upon discharge ELM suffered from a residual nominal aphasia and mild memory impairment that later disappeared. In August of 1985 he was readmitted to the Montreal Neurological Hospital, presenting with pronounced anomia, memory impairment, and dysgraphia. A CT scan conducted on August 9 revealed irregular enhancing lesions deep in the left mesiotemporal lobe (areas 21 and 37) and deep right mesiotemporal lobe. His condition improved and he was discharged on August 21, 1985.

Neuropsychological Assessment

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

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

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

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

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

In terms of face recognition ELM can match to sample a target face embedded in a row of distractor faces (32/32 trials). He cannot, however, discriminate familiar from unfamiliar faces or previously viewed from novel unfamiliar faces. He cannot identify emotional expressions (13/42). On a corpus of 48 faces whose identities were well known to him, ELM failed to identify a single face.

METHODS

Experiment 1

Procedure

Faces of male models taken from fashion magazines were digitized and presented on a computer screen. Digitized, gray-scale faces were 9 cm high by 6 cm wide. Eight such faces were divided into two sets with each set being informally matched for visual similarity. The four faces of one set were randomly paired with four hockey players' names, and the other four faces were randomly paired with four names of people with different occupations. Face-name pairings remained consistent throughout experimentation.

Prior to testing, ELM was asked to give semantic information about the following eight names; Kirk Muller (forward for the Montréal Canadiens Hockey Club), Vincent Damphousse (Canadiens' forward), Benoit Brunet (Canadiens' forward), Mark Recchi (Canadiens' forward), Matthew Broderick (actor), Jean Charest (politician), Larry Walker (baseball player), and Jacques Villeneuve (racecar driver). He was able to do so. ELM was then seated in front of a computer screen and told that he would see a number of faces accompanied by the names of hockey players. He was given a list of the relevant names to which he could refer.

On learning trials, faces were presented one at a time, and ELM was allowed to view them for as long as he wanted. He then clicked the mouse, which initiated a digitized recording of one of the names. On test trials ELM was shown the face and asked to generate the name. He was given as much time as necessary to generate these names. For both learning and test trials, faces were presented in one of three orientations: vertically or rotated 25° to the left or right of vertical in the picture plane. This manipulation was used to prevent ELM from using nonface local cues (e.g., chin pointing to the corner of the screen) to "recognize" the face stimuli.

Eight learning trials (two of each face-name pairing) were followed by eight test trials. Trials were presented in random order with the proviso that no two identical faces followed one another. This pattern of 8 learning trials and 8 test trials was repeated 12 times, yielding 96 learning and 96 test trials per block. Three such blocks were administered with a 10-min break between blocks.

All stimuli were presented using Psychlab software (Bub & Gum, 1990). This application was run on a Power Macintosh connected to an Apple Color Plus monitor.

Experiment 2

Control Subjects

Two sets of six elderly control subjects were tested. (Mean ages = 71.16 for Set 1 and 69.67 for Set 2.) All controls were living independently, either alone or with a spouse, and had no subjective complaints concerning either memory or face recognition.

Faces

Eight synthetic faces were constructed using a template female face. Faces were constructed using one of two eye, nose, or mouth sizes. The eyes, nose, and mouth were removed from the template face and replaced with unique combinations of these three features. Faces were constructed using Adobe Photoshop. Faces were divided into two sets of four faces each. In one set the faces all had large noses. Face 1 had small eyes and a small mouth (1,1), Face 2 had small eyes and a large mouth (1,2), Face 3 had large eyes and a small mouth (2,1), and Face 4 had large eyes and a large mouth (2,2). The faces in the second set all had small noses. The four faces of set 2 used the same combinations of eye and mouth sizes as in set 1.

An additional three faces were distorted using Adobe Photoshop (twirl, pinch, and spherize effects).

For ELM, and one set of controls, Face (1,2) was removed from the plausible face set, and the spherically distorted face was removed from the distorted face set. Testing was conducted using two triads of plausible faces and one triad containing distorted faces.

Prior to experimentation, subjects were asked to give semantic information about the following six names; Tonya Harding (former ice skater), Nancy Kerrigan (ice skater), Josée Chouinard (ice skater), Hillary Clinton (First Lady), Celine Dion (singer), and Betty Grable (actress). All subjects were able to do so.

The two sets of faces were then paired to semantically close and disparate names using an ABBA design. Face-name assignments are shown below the relevant conditions in Figure 3. Note that exact same faces (A_1 and B_2, B_1 and A_2) were mapped to semantically close and disparate sets of names in separate conditions.

Subjects were then seated in front of a computer screen and told that they would see a number of faces accompanied by names that were familiar to them. Subjects were given lists of the names relevant to that session's testing, for reference.

On learning trials, faces were presented simultaneously with a digitized recording of one of the names. Faces remained on screen for 2200 msec followed by a blank screen. After an intertrial interval of 1500 msec, a new face-name combination would appear. Six learning trials were presented in random order with the proviso that no two faces followed one another and that there were two repetitions of each face-name pairing.

Six learning trials were followed by six test trials. On test trials a "ready" prompt was given immediately followed by one of the faces presented in learning trials. Subjects were given as much time as necessary to generate the name that was associated with the face. Faces were presented in random order with the proviso that no two faces followed one another and that there were two presentations of each face within a six test trial sequence.

Six learning trials (two of each face-name pairing) were followed by six test trials. This pattern of 6 learning trials and 6 test trials was repeated 24 times, yielding 144 learning and 144 test trials. All faces were presented vertically. A 10-min break was given after the seventy-second test trial.

For ELM, each cell of the ABBA design was tested on a different day. For set B_2 after having made four errors early during testing, ELM demonstrated flawless face-name matching (41 correct in a row). Testing on this set was therefore discontinued after 72 test trials. For controls, sessions A_1 and B_1 were tested in a single session, and A_2, B_2 and the distorted condition were tested in a separate session. For controls perfect identification was assumed after 12 consecutive correct test trial responses, whereupon testing for this set was discontinued.

For all subjects the triad containing the distorted faces was presented last. These faces appear as the "Distorted Skaters" set in Figure 3.

For the second set of control subjects all eight plausible synthetic faces were used, along with a quadruplet of faces containing one plausible and three distorted faces. To accommodate these extra faces, two extra names were added (ice skater Katerina Witt and tennis player Steffi Graf) to the sets of semantically related and unrelated names. Eight learning trials (two of each face-name pairing) were followed by eight test trials (in which faces were presented and subjects asked to "name" them. This pattern was repeating until subjects correctly named 16 consecutive test trial faces, whereupon perfect identification was assumed, and testing for that set discontinued. Sets A_1 and B_1 were tested in one session, and Sets B_2, A_2 and the distorted condition were tested in a separate session on another day.

Acknowledgments

The authors would like to express our heartfelt appreciation to our colleague and friend ELM without whose unending patience and altruism this research could not have been conducted. We would like also to thank the subjects who served as healthy controls in Experiment 2. The first author received support from a post-doctoral fellowship from the Alzheimer Society of Canada and supplemental funding from the EJLB Foundation. This research was also supported by grants from the Medical Research Council of Canada to the second and third authors. Martin Arguin is a Chercheur-Boursier of the Fonds de la Recherche en Santé du Québec.

Reprint requests should be sent to Mike Dixon, Dept. of Psychology, University of Waterloo, Waterloo, Ontario, N2L 3G1, Canada, or via e-mail: mjdixon@watarts.uwaterloo.ca.

Notes

1. Initially eight control subjects were tested in the ABBA design. As a group these eight controls made significantly more errors for the skaters relative to the unrelated sets. Two control subjects were discarded because of marked proactive interference effects. Within the *A* and *B* conditions of the $A_1B_1B_2A_2$ design, these two subjects showed marked interference from A_1 to A_2 (16 to 47 errors and 40 to 47 errors, respectively) and especially from the B_1 to B_2 conditions (5 to 55 errors and 29 to 53 errors, respectively). Because ELM and the other subjects showed either constant or, most often, improved performance (i.e., practice but not interference effects), these two subjects were deemed inappropriate controls for ELM on this paradigm and were not included in the reported analyses.

2. We conclude that these findings reflect an interaction of visual and semantic proximity on *identification* performance rather than simply on memory performance for paradigms in which stimuli and required responses overlap. We base this conclusion on two lines of evidence. First, in the experiments designed to unravel ELM's category-specific visual agnosia for objects, we used the exact same paradigm as in the present study except that blobs were used instead of faces and object names were used instead of person names. In this study there was an overwhelming correspondence between his performance with conjunction sets of blobs and his real life identification problems for objects with overlapping visual and semantic features (e.g., conjunction sets of blobs mapped to semantically close concepts such as birds, cars, animals, or stringed musical instruments could not be identified in the laboratory paradigm nor could real birds, cars, animals, and stringed musical instruments be identified in real life).

Secondly, if blobs forming a conjunction set are presented in different locations on a computer screen and ELM is asked to remember the locations of these blobs, he performs quite poorly (50% after an initial learning period; Dixon et al., 1997). If the input stimulus becomes conjunction sets of sounds (e.g., unique combinations of pitch 250 versus 2500 Hz) and reverberation (sine wave versus square wave modulation) presented on speakers placed at different locations (comparable to the four computer screen locations), ELM performs flawlessly on this task (0% errors after an initial learning period). This finding is not surprising given that ELM has category-specific visual agnosia, but what is of importance is that on each of the above tasks the required response on test trials was a location, and the input stimuli were conjunction sets of sensory stimuli. Thus, two paradigms in which there was substantial overlap between stimulus and response sets yielded marked performance differences. This suggests that ELM's performance on laboratory tests of the kind used in the current study reflect true identification problems for visually and semantically close stimuli rather than simply problems performing on paradigms in which there is overlap between stimulus and response sets.

REFERENCES

Arguin, M., Bub, D., Dixon, M., Caille, S., & Fontaine, S. (1996). Shape integration and semantic proximity effects in visual agnosia for biological objects: A replication. *Brain and Cognition. 32,* 259–261.

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

Beymer, D., & Poggio, T. (1996). Image representations for visual learning. *Science, 272,* 1905–1909.

Bub, D., & Gum, T. (1990). Psychlab software. Montreal: McGill University.

Burton, A. M., Bruce, V., & Johnston, R. A. (1990). Understanding face recognition with an interactive activation model. *British Journal of Psychology, 81,* 361–380.

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

De Renzi, E., Faglioni, P., Grossi, D., & Nichelli, P. (1991). Apperceptive and associative forms of prosopagnosia. *Cortex, 27,* 213–221.

Dixon, M., Bub, D. N., & Arguin, M. (1998). The interaction of object form and object meaning in the identification performance of a patient with category specific visual agnosia. *Cognitive Neuropsychology,* (in press).

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

Farah, M. J. (1991). Patterns of co-occurrence among the associative agnosias: Implications for visual object representation. *Cognitive Neuropsychology, 8,* 1–19.

Farah, M. J., O'Reilly, R. C., & Vecera, S. P. (1993). Dissociated overt and covert recognition as an emergent property of a lesioned neural network. *Psychological Review, 100,* 571–588.

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

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

Lissauer, H. (1890). Ein fall von seelenblindheit nebst einem beitrage zur theori derselben. *Archiv für Psychiatrie und Nervenkrankheiten, 21,* 222–270.

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

McNeil, J. E., & Warrington, E. K. (1991). Prosopagnosia: A reclassification. *Quarterly Journal of Experimental Psychology A: Human Experimental Psychology, 43,* 267–287.

Newcombe, F., Young, A. W., & De Haan (1989). Prosopagnosia and object agnosia without covert recognition. *Neuropsychologia, 27,* 179–191.

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

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

Schweich, M., & Bruyer, R. (1993). Heterogeneity in the cognitive manifestations of prosopagnosia: The study of a group of single cases. *Cognitive Neuropsychology, 10,* 529–547.

Sergent, J., & Poncet, M. (1990). From covert to overt recognition of faces in a prosopagnosic patient. *Brain, 113,* 989–1004.

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: Further fractionations and an attempted integration. *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.
Young, A. W. (1994). Covert recognition. In M. J. Farah & G. Ratcliff (Eds.), *The neuropsychology of higher level vision* (pp. 331–358). Hillsdale, NJ: Erlbaum.
Young, A. W., Hellawell, D., & De Haan, E. H. F. (1988). Cross-domain semantic priming in normal subjects and a prosopagnosic patient. *Quarterly Journal of Experimental Psychology A: Human Experimental Psychology, 40,* 561–580.